Restraining Orders and the Deterrence of Domestic Violence*

Eric Helland†
Claremont McKenna College and RAND

Minjae Yun‡
Claremont McKenna College

December 15, 2025

Abstract

This study evaluates the causal impact of civil restraining orders on crime, with a focus on domestic violence cases in Los Angeles from 2020 to 2023. Using a dataset of all restraining order requests in Los Angeles and matched LAPD crime data, we employ an event study and a judge leniency/examiner instrumental variable (IV) design. Results from the event study indicate a decrease in crime likelihood following the issuance of the order. Our IV analysis reveals that granting a restraining order reduces crime probability by 33%-41% in the subsequent year, with significant deterrence effects observed primarily in domestic violence and aggravated assault-related cases, but only for domestic violence restraining orders with minor children present. The results from domestic violence restraining orders with minor children present, but not other types of restraining orders, are consistent with two overlapping theories. The first is that domestic violence restraining orders remove legal or financially required interactions between petitioners and respondents that are nearly unavoidable without an order when the petitioner and the respondent have minor children. The second is Aizer and Bo's theory of time-inconsistent preferences, in which restraining orders represent a commitment device. The study contributes to the literature by providing causal estimates of the deterrence effects of restraining orders and insights into the mechanisms by which restraining orders have an impact.

*The authors wish to thank Paul Heaton, Aurélie Ouss, Jennifer Doleac, Aviv Caspi, Yotem Shem-Tov, and seminar participants at ALEA (NYU) 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).

‡Lowe Institute of Political Economy, Claremont McKenna College, 888 Columbia Avenue, Claremont, CA 91711 (minjae.yun@cmc.edu).

1 Introduction

The CDC estimates that the lifetime prevalence of domestic and intimate partner violence (IPV), including physical violence, sexual violence, and stalking by an intimate partner, is more than 1 in 3 for women (35.6%) and more than 1 in 4 men (28.5%) (Black et al., 2011). Moreover, victims often suffer repeated abuse by their partner (Aizer and Dal Bo, 2009). In Los Angeles County, the location of our study, a 2010 survey by Blue Shield of California found that IPV was similarly prevalent, with one in five respondents having experienced domestic violence at least once in their lifetime (see Los Angeles County Department of Public Health, Injury and Violence Prevention Program (2013)). The consequence of this violence is often fatal. In Los Angeles County, a woman is killed by an intimate partner every five days (Tjaden, 2000).[1] In the 1970s, in an attempt to deal with the inadequacies of the legal remedies available to victims, states began to enact three different policy reforms. The first was implementing mandatory arrests or policies encouraging officers to arrest individuals accused of IPV. The second, so-called no-drop policy required prosecutors to continue to prosecute individuals charged with domestic violence even if the victim stopped cooperating with police and prosecutors (see Aizer and Dal Bo (2009)). The policies often coincided with the creation of specialized units to address IPV in prosecutors' offices. The third major reform was adopting civil restraining orders, which gave abuse victims a civil, rather than criminal, remedy to prevent contact with the perpetrator.[2] Civil restraining orders are designed to offer prompt protection to victims by legally preventing abusers from further violence. In most states, including California, violating a restraining order is considered a criminal offense, allowing law enforcement to arrest and potentially convict the abuser for non-compliance. However, as discussed below, the prevalence of such arrests without other charges is rare.[3] Given the widespread occurrence of IPV, it is unsurprising that a substantial body of literature addresses its consequences and prevention. Much of the research on IPV has focused on the changes to policing methods adopted after the 1970s and how these methods impacted the likelihood of IPV.[4]

By contrast, relatively little research has focused on the second major intervention, civil restraining (or

[1] Los Angeles is not an outlier; according to the Bureau of Justice Statistics (Truman and Morgan, 2014), intimate partner violence amounts to 15 percent of all violent crimes.

[2] In 1976, only two states had laws allowing civil restraining orders. By 1994, all 50 states had some form of protective orders, often with specialized courts dealing with the orders (Ko, 2001).

[3] Civil protection orders are, in some ways, an alternative to criminal charges due to their immediate availability to eligible victims and broad protection scope. Moreover, criminal protective orders are only available after an arrest and are often more limited in scope (see Ko (2001) and Judicial Council of California (2023)).

[4] For example, one pro-arrest policy, warrantless domestic violence arrests, which either requires police to arrest suspects regardless of the wishes of the victim (mandatory arrest) or provides officers with a reduced threshold for arrest, has been extensively studied. The mandatory arrest intervention was first examined in a series of randomized studies on the effectiveness of arrests on domestic violence (see Sherman and Berk (2017), Dunford et al. (1990), Hirschel and Hutchison iII (1992), Sherman et al. (1992), Pate and Hamilton (1992), and Berk et al. (1992)). These studies produced inconclusive findings mainly due to non-compliance with random assignment. A series of papers have used difference-in-difference, event studies, or intent-to-treat (ITT) instrumental variables to examine the impact of this policy. In the earliest of these studies, Iyengar (2009) finds evidence that mandatory arrests increase violence. A replication by Chin and Cunningham (2019) examines Iyengar's study and finds no evidence that mandatory arrest policies increase spousal homicides, and finds some evidence that discretionary arrest laws reduce spousal homicides. Finally, using a leniency design similar to this study, Amaral et al. (2023) find evidence that IPV arrests reduce future IPV. The second major reform, no-drop prosecution policies, is investigated by Aizer and Dal Bo (2009).

protective) orders, despite the fact that restraining orders are arguably the primary form of protection for IPV victims (Ko, 2001). A series of studies in the 1990s (see Buzawa and Buzawa (1996)) examined restraining orders.[5] For example, Ko (2001) documents five studies between the mid-1980s and 2000 addressing the reduction in physical abuse from civil restraining orders. Most of these studies were interviews with women in domestic violence shelters (see Chaudhuri and Daly (1992)). However, two studies use arrest data. The first is a study of 663 restraining order cases in Quincy, Massachusetts (Klein, 1996). Klein and coauthors found no significant differences in re-abuse rates for those who dropped or kept their restraining orders. The second, Carlson et al. (1999), studied court and police records and found a sizable decrease in arrests relative to the previous two years for domestic violence after orders were issued. Moreover, permanent restraining orders had a larger effect than temporary restraining orders (TROs).

The problem is that all studies of the effectiveness of civil restraining orders suffer from various methodological and generalizability issues that make inference difficult. For example, a control group is absent in all of the studies. Moreover, these studies were conducted in the mid-1980s and 1990s, when reforms were relatively new, and civil courts and law enforcement were still developing procedures to deal with civil restraining orders. Over the following decades, courts and law enforcement developed procedures to streamline the process (see our discussion below), raising concerns about the applicability of these early studies. Perhaps more problematic, even for the studies by Klein (1996) and Carlson et al., which utilize arrest data, is that requesting and granting the order is likely endogenous. Victims who are granted the order may differ greatly from those who are not. Our task in this paper is to estimate the causal effects of the court granting the petitioner's request for a restraining order on the likelihood of a crime, given the obvious endogeneity created by the fact that judges are granting orders based on the likelihood the respondent will be violent in the future, i.e., commit a crime against the petitioner.

Our primary data source consists of information on 20,384 restraining order requests filed in the city of Los Angeles between 2020 and 2024, which were scraped from the Los Angeles Superior Court's online case access system for family court and which we are able to match to block-level address data from the LASC.[6] Our sample consists of 20,384 requests for orders, of which 4,803 were granted. We use the 15,328 orders that were not granted restraining orders as our control group.[7]

The majority of the order requests have a temporary restraining order in place until the hearing. Thus, the

[5]Studies of the physiological benefits of restraining orders are more common, although these studies also suffer from similar methodological problems. (see Ko (2001)).

[6]From the full sample, we randomly selected 4,000 requests for a restraining order and hand-collected the addresses covered by the order request. We find 3,532 orders with a specific address from which the respondent to the order, the alleged abuser, would be barred if the order were approved, consistent with the match rate to LASC addresses. We compare the addresses at the block level to the addresses provided by LASC and find that, for orders with an address, only about 1.5% are different between the two samples.

[7]About 20% of orders in our sample do not have a home address that the petitioner requests the respondent be barred from. Typically, these are civil orders concerning stalking that occur away from the petitioner's home and seek a "stay-away" provision, barring the respondent from contacting the petitioner. These orders typically leave out the petitioner's home address for obvious reasons. Given that the nature of estimation requires us to observe crimes at a specific block, we remove these observations from the sample.

treatment in our study involves transitioning from having a restraining order to continuing that order after the judge grants a 3-5 year order at the hearing, or from having an order to not having an order if the judge denies the order request at the hearing.

We then match the address in the order request to a database of crimes in the city of Los Angeles. Like the order data, crime data is only available at the block level. Thus, the unit of analysis is the probability that a residential domestic violence crime was reported in the block of the address found in the petition during the year after the hearing at which the restraining order was granted or denied.

We use two estimation techniques to determine the effect of granting a civil restraining order request on the probability of crime. First, in our preliminary data analysis, we estimate an event study post-21-day hearing, comparing petitions with minor children in the household that are granted a domestic violence restraining order at the hearing (ROAH) to those requests denied by a judge. We examine the daily probability of a crime in the year following the hearing. Because the event study does not control for the fact that orders receiving a ROAH may differ significantly from those in which the judge denies the ROAH, our primary estimation is a leniency design IV based on the random assignment of judges and the likelihood that the assigned judge will grant a ROAH.

In our event study model, which includes petition (case)-fixed effects, we find that a granted domestic violence restraining order petition with minor children present is associated with a higher probability of crime per day over the next year at the block level but that this probability declines in magnitude at an increasing rate over the year following the granting of the order relative to petitions rejected by a judge.[8]

When we estimate the probability of a crime in the subsequent year at the block-level address without controlling for endogeneity, we again find that granting a restraining order is generally associated with an increase in the likelihood of a crime. This suggests that the likelihood of getting an order from the court is not random and that the court is granting orders in cases where the threat of a crime is higher.

Our second estimation strategy leverages the as-if random assignment of judges to preside over the mandatory 21-day hearing. Because judges differ systematically in their propensity to grant restraining orders, this variation identifies a local average treatment effect (LATE): the causal effect of granting a restraining order for cases whose outcomes are influenced by judicial assignment, independent of the selection biases present in our event study and OLS estimates (see Agan et al. (2023); Chyn et al. (2025)). When we control for endogeneity using our leniency design IV, we find that granting an order is associated with a 26%-40% percentage point drop in the probability of a residential crime meeting our definition reported on the block of the address on the order request, but only for domestic violence restraining orders with minor children present.

This is surprising, as the penalties for violating a restraining order are the same across all order types. Hence,

[8]Note that we have no blocks in which we have two restraining order requests so petition fixed effects are equivalent to block fixed effects.

the deterrence effect of granting an order should be similar.[9] There are, however, several possible theories explaining the finding that civil restraining orders' deterrence effect is primarily on domestic violence crimes. The first is that when a judge denies the petitioner's request for an order in a petition that is not a DVRO with minor children present, the petitioner has other possibilities to reduce the likelihood of being a victim of a crime perpetrated by the respondent. For example, they could move, avoid contact, or generally make it difficult for the respondent to find them. By contrast, absent a restraining order, petitioners who share custody of minor children or housing with the respondent, or who require spousal support, are unlikely to be able to avoid contact with the respondent, as they are more likely to be legally or financially required to interact with the respondent without the restraining order. The second theory that is consistent with our findings is Aizer and Dal Bo (2009)'s theory that the high rate at which domestic violence victims return to their abusers can be understood as a form of time-inconsistent preferences.[10] Our event study results are consistent with Aizer and Bo's theory in that their model predicts that as the temporal distance from the violence that precipitated the petitioner's request for a restraining order grows, additional violence would become more likely as the probability that the petitioner allows the respondent to return, but again only for DVROs with minor children present. In this theory, the DVRO is a commitment device. After a permanent restraining order is granted at the 21-day hearing, it is difficult for the petitioner to rescind it. They must request the removal of the court order and secure the judge's agreement. Thus, the cost of relenting and allowing their abuser to return is much higher once the order is in place, and, perhaps more importantly, the costs for the abuser to push the victim to allow them to return are higher with an order in place.

To move beyond our LATE estimates, we estimate marginal treatment effects (MTEs), which trace how the causal impact of granting a DVRO varies with unobserved case characteristics that make a grant more or less likely (Heckman and Vytlacil 2007a,b). The results show little evidence of treatment-effect heterogeneity: the deterrent impact of granting a DVRO appears broadly consistent across cases, implying that raising the overall grant rate would still reduce violence, though possibly with somewhat smaller effects than our LATE estimate.

This paper contributes to the literature on civil restraining orders in two ways. First, we provide causal estimates of the impact of granting a restraining order on the likelihood of a crime. Second, our paper provides some evidence of how civil restraining orders reduce the likelihood of crime, specifically by providing a legal mechanism for petitioners with domestic relationships to avoid contact with the respondent and/or as a commitment device for petitioners (abuse victims) to maintain separation.

[9] 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 (criminal offense under California Penal Code § 273.6), and both have similar aggravating factors that can escalate the offense to a felony with state prison time of 3 years or higher and fines. Theoretically, the value of the crime to the perpetrator could be greater for respondents covered by civil harassment orders, making them harder to deter, but that seems contrary to the conventional wisdom that domestic violence abusers are harder to deter.

[10] In Aizer and Bo's model, the high degree of cyclicality in domestic violence results from time-inconsistent preferences. The victim's preferences change as the initial violence grows more remote, changing the relative value of the relationship to the disutility of the potential future violence.

The paper proceeds as follows: Section 2 provides background on the legal framework for restraining orders and how this framework informs our research design. Section 3 discusses the data for our study. Section 4 presents our research design and the results from an event study. Section 5 presents the results from a leniency design instrumental variables regression, including robustness checks, and Section 6 concludes.

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.[11] 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)).[12] Violations of civil restraining orders can also carry a misdemeanor or felony charge.[13]

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

[11] 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.

[12] 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.

[13] 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).)

whom the petitioner has or once had a relationship, such as a spouse, partner, or close relative.[14] 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 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.[15]

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 temporary 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.[16] 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

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

[15] 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.

[16] 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.

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). 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.[17] 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.[18] 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.[19]

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

[17]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.

[18]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.

[19]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.

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.[20]

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.[21]

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.[22]

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

[20] One concern is that petitioner's failure to appear may be driven by court logistics that could undermine randomization. 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. We also estimate a CJIVE specification that uses judge availability on a given day.

[21] 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.

[22] 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.

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 our 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 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.[23] 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 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.[24]

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.[25] 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

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

[24] 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.

[25] 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.

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.[26] 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.[27] 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.[28] 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.[29]

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.

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

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

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

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

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.[30] 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 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.[31]

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.

[30] 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.

[31] 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.

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 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.[32]

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.

Our first approach to estimating the model is an event study model that estimates the daily probability of a residential crime on the block of the address in the restraining order request. The event study model compares the likelihood of a reported crime before and after the restraining order is granted to the similar probability for those requests that were denied at the hearing stage. This approach does not directly address the issue of under-reporting being correlated with restraining order status or endogenous selection unless you can reasonably assume that the propensity to report crimes is unaffected by whether the restraining order is granted. It seems reasonable that this would lead to under-reporting among denied petitioners, biasing the results toward zero.

Our second approach uses 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,

[32]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%).

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

As noted above, we use two different approaches to estimating the impact of the court granting a civil restraining order on crime. As preliminary data analysis, we estimate an event study model in which we compare the daily probability of a crime at the block level address for restraining order requests granted a restraining order relative to requests denied at the 21-day hearing. For this analysis, we focus only on DVROs in which the petitioner and respondent share minor children.[33]

Our second approach addresses the endogeneity of an order being granted and the probability of a crime by utilizing 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.

4.1 Preliminary Analysis: An Event Study Approach

Because we know the exact date of the RO hearing and because the decision at the 21-day hearing represents either a continuation of a restraining order or a change in status to no longer have one, we treat the RO date as our event. We also know the exact date of the crime at the block level, allowing us to estimate the daily probability of a crime at the block level if the restraining order is granted, relative to those requests that were denied. We use our measure of residential crimes labeled by the LAPD as domestic violence as our outcome measure.

Specifically, we estimate the equation,

$$Pr(Crime_{it}) = \beta_1 grant_{it} + \sum_{k \neq 3} \beta_k \text{bins}_k + \sum_{k \neq 3} \beta_k \text{bins}_k \cdot \text{grant}_{it} + \alpha_i + \gamma_t + \varepsilon_{it}, \quad (1)$$

where $Pr(Crime_{it})$ an indicator variable equal to one of a crime was committed on day t within the block of the address contained in restraining order i, bins_k represents the time-bins for week 2, 3, 4, 5, weeks 5 through 10, 10-20 and greater than 20 weeks, and $grant_{it}$ equals one if the restraining order after hearing was granted and zero if the order was denied at the 21-day hearing. The model also includes case fixed effects, α_i, and year-month fixed effects, γ_t; the standard errors are clustered at both the order request and

[33] We find no effect over time when we aggregate restraining orders across different types.

date to allow for arbitrary correlation in the errors within each cluster. Because the order almost always occurs at 21 days, the reference period is period 3.

We use seven weekly event indicators rather than daily event estimates to reduce the daily noise and improve precision. This does not inherently create problems in an event study model as long as the binning accurately reflects the timing of treatment and the outcomes within each bin. However, the interpretation shifts to average effects over weekly rather than daily changes.

Figure 2 and Table A1 present the results for DVROs with minor children. The 21-day (week 3) hearing date is the reference point, with the line at zero. The table includes the three periods (week 1 $\times$ grant, week 2 $\times$ grant, and week 3 $\times$ grant) covered by the temporary order, with week 3 again as the reference point. The coefficient on granting, which is not shown in the figures, is positive and significant across all of our crime measures. Taken literally, this would suggest that granting a restraining order causes crime. In light of our findings below, it seems more likely that the likelihood of a restraining order being granted by the judge is endogenous. That is, judges are granting orders in cases where they perceive the threat of a crime is higher, and judges are, on average, correct in this assessment. In the next section, we turn to our IV strategy to address this endogeneity. However, in our preliminary data analysis, we ignore this endogeneity and focus on how the probability of a crime evolves after the 21-day hearing.

In Figure 2a, we plot the weekly bin for cases where the restraining order was granted. After week 4 or 5, depending on the crime measure, we observe a statistically significant drop in the daily probability of a crime that persists for the subsequent year. We find a similar pattern in Figure 2b, in which we estimate the model using assaults. We also see a similar pattern for aggravated assaults (Figure 2c).

One interesting feature of our results is that the decline in the daily probability of crime grows over time for granted orders relative to denied orders. This is consistent with Aizer and Dal Bo (2009)'s theory that the high likelihood that victims of domestic violence return to their abusers is due to time-inconsistent preferences. In their theory, the victim's preferences change with time for the violent incident that precipitated the request for a restraining order. Right after the incident, the victim's relative value of the relationship is low, but it increases over time. Thus, without the restraining order, we would expect the probability of crime to grow over time as victims relented and allowed the abuser back into their lives. In this theory, because once the DVRO is granted, it is relatively difficult for the petitioner to get the order rescinded, the 21-day hearing serves as a commitment device for the petitioner, as it raises the cost to the abuser of renewing the relationship.

4.2 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, rather than the petition–time structure used in the event study.

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

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 or CJIVE method (described below). $\mathbf{X'}_i$ 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.[34] We employ the cluster-jackknife instrumental variables estimator (CJIVE) in addition to the residualized JIVE. As noted by Frandsen et al. (2025), CJIVE remains unbiased in the presence of within-shift, in our context, calendar day, correlations. In other words, bias may arise not only from the current judge but from the entire cluster of cases within the same shift. Excluding the full shift, therefore, helps eliminate this potential bias. Similar to the setting studied in Frandsen et al. (2025), the family court also operates at the daily shift level. We construct CJIVE using the full roster of judges across courts, since judges who don't hear restraining orders on a given day may nevertheless have been available during the "shift."[35]

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.

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

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

[35] In practice, it is rare for a judge not assigned to the standalone restraining order courts to hear even a single restraining order case on a given day.

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 coefficients are smaller, .502, .469, and .408 for the CJIVE estimator, suggesting we may have a weak instrument problem. The first stage F statistics are 146.587, 144.845, and 134.046 for the residualized JIVE estimator and 74.299, 64.59 and 57.762 all far in excess of the value of 10 used to identify weak instruments.[36]

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 3.

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, such an effect would bias our analysis against finding any impact of protective orders, thereby mitigating

[36] 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.

concerns about overestimating their effects. Moreover, our estimates are based on reported crimes rather than arrests, partially 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 3 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 3. 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 'average elasticity,' 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) (C-OLS) reweights the sample to match the complier distribution implied by our first stage. Comparing Columns (2) and (3), we find that the compliance-weighted OLS estimates are closer to zero, though not statistically distinct from the unweighted OLS estimates.

In columns (4)-(6), we control for endogeneity using our residualized JIVE instrument constructed using the previous two years of data, discussed in the previous section. The coefficients are larger in magnitude, although it is not significant, indicating a 7.44 percentage point reduction in the probability of a domestic violence crime at the block-level address. The magnitude of this drop produces an elasticity of -.45, indicating a 45% decrease in the likelihood of a crime over the year after the order is granted. When we include control variables, we find that the impact is similar, a 6.69 percentage point reduction in the probability of a crime. When we construct the residualized JIVE instrument using all previous years of data, our results are not appreciably more precise and remain insignificant.

In column (7)-(9), we use our CJIVE measure. While residualized JIVE partials out observed calendar and case-type effects, CJIVE further eliminates potential bias from unobserved shocks shared by all cases on the same daily calendar, yielding an instrument that is more robust to within-day correlation. The estimates are positive but continue to be insignificant.

In Panel B, we reestimate the model with assaults rather than police identified domestic violence. The OLS estimates are negative with and without control variables, and the results are similar for our residualized JIVE and CJIVE estimators. In both cases, the impact of granting a restraining order after the 21-day hearing is insignificant.

We find a somewhat different pattern in Panel C, which focuses on aggravated assault cases. Both the OLS and C-OLS estimates are positive but statistically insignificant, with the C-OLS coefficient smaller in magnitude. This pattern suggests that the association between restraining orders and aggravated assaults is weaker among compliers, indicating that the complier subset may be somewhat less correlated with elevated assault risk ex ante. Similarly, the JIVE and CJIVE estimators are insignificant. In Panel D, we estimate the model by including only simple assault. As with all assaults, the impact of granting a restraining order is insignificant both with and without instruments.

Although the results are statistically insignificant across the board, they nonetheless provide insight into the nature of selection in judicial decision-making at the 21-day hearing. Bjerk (2025) argues that the gap between OLS and IV coefficients can only arise from selection into treatment or heterogeneous treatment effects. Consistent with this logic, the IV estimates are significantly larger in magnitude and negative in

sign compared to the OLS estimates. If the difference is due to selection, then the OLS estimates reflect both the causal effect of granting a restraining order and selection bias arising from unobserved differences in petitioners' risk of subsequent crime. Judges are more likely to grant orders to petitioners who present stronger evidence of threat, meaning that those who receive a restraining order are also more likely, ex ante, to be victims of crime. If instead heterogeneity drives the difference, it would mean that compliers (those whose treatment status is determined by judge leniency) experience different causal effects than always-treated petitioners. Given the statistically indistinguishable OLS and compliance-weighted OLS estimates, we find little evidence that compliers differ systematically from the overall treated population in observable characteristics or treatment responsiveness. This suggests that the OLS–IV gap arises primarily from selection into treatment on unobservables—judges are more likely to grant orders in high-risk cases, rather than from heterogeneity in causal effects across compliers (see Bjerk, 2025, Agan et al., 2023).

5.1 Breakdown by Type of Order Requested

In Table 5, we estimate the model of four subsamples of restraining order requests using police identified domestic violence crimes, our preferred measure of crime. As shown in Table 2, the block-level crime rate is similar across the different order types. The first is civil harassment orders, which, as discussed above, are filed against people with whom the petitioner has no close relationship. The coefficient is statistically insignificant for both the OLS and IV estimates, and the estimated impact is positive.

In Panel B, we estimate the model using only DVRO, restraining order requests against spouses, former spouses, close relatives, or other intimate partners. Here, we find small negative and insignificant effects in the OLS estimation. We find that the magnitude grows once we correct for endogeneity. For the IV estimates, we find that negative but again insignificant effects are similar to our estimates based on pooling restraining orders. The impact ranges from a 2.12 to 8.95 percentage point drop in the probability of a crime in the block of the address over the next year.

In Panel C, we estimate the model on DVRO petitions in which minor children are in the household. The OLS results are positive, but in the case of DVROs with minor children, we find significant and much larger effects ranging from a 33.7 to 41.5 percentage point decline in the probability of a police defined domestic violence crime, which translates into a 184 to 227% decline. The impact is significant whether we use the residualized JIVE or CJIVE. Whether we include control variables or not, and when we use the Anderson-Rubin correction, the results are statistically significant.

In Panel D, we examine only DVROs in which minor children are not present. None of the estimates are statistically significant, and the magnitude is considerably smaller than the estimates of DVROs with minor children.

In Table 6, we estimate the model using our different measures of crime at the block level, and we restrict the sample to requests for a DVRO in which the petitioner has minor children. Panel A repeats Panel C

from Table 5 for reference. For all assaults, Panel B, we find a negative and statistically significant impact of granting the order on the probability of an assault over the next year, only for our CJIVE estimator constructed with the previous two years of data. Moreover, the AR confidence intervals all include zero. All other specifications are insignificant.

In Panel C, we estimated the model using the probability of a residential aggregated assault at the block level in the year after the 21-day hearing. We find a negative and significant impact for our CJIVE estimator using the current and last years to construct the instrument. The impact is -19.9 and -19.5 percentage points and translates into a larger elasticity than the estimates for police defined domestic violence, 334.8% and 327.5% reduction in the probability of a crime. Essentially, granting a DVRO when minor children are involved reduces the probability of an aggravated assault at the block listed on the order request almost to zero in the subsequent year. While residential aggravated assaults occur infrequently in our sample, they are among the most severe offenses, involving weapons or serious injury, so deterring these incidents may generate disproportionately large social benefits. In Panel D, we find insignificant and both small positive and negative deterrence effects for simple assault.

One interesting feature of the results is that the impact of restraining orders appears to be confined to domestic violence crimes and aggravated assaults at residences, and only for domestic violence restraining orders petitions with minor children present. Before turning to a possible mechanism for the pronounced effect of domestic violence restraining orders in cases with minor children, we examine several robustness checks restricting our sample to DVROs with minor children.

5.2 Robustness Check: Alternative Estimation Methods

One concern with our sample period (January 1, 2020, to December 31, 2023) is that its early months coincide with the COVID-19 lockdowns in Los Angeles. During the pandemic, the Los Angeles Superior Court treated restraining-order proceedings as essential and continued hearing them remotely, primarily by video or telephone, from March 15 to June 5, 2020, when the stand-alone restraining-order courts at Stanley Mosk resumed in-person hearings.[37] To ensure that virtual hearing procedures do not influence our estimates, Table 7 re-estimates the model excluding petitions filed during this remote period. The results are nearly identical to those using the full sample: the IV estimates (both rJIVE and CJIVE) remain statistically significant and indicate that granting a petition reduces the likelihood of a police-identified residential domestic-violence crime at the petitioner's address block by approximately –0.33 to –0.41.

In Table 8, we assess the robustness of our main results to alternative ways of constructing the instrumental variable. In Columns (1)–(2), we estimate the model using 2SLS with judge fixed effects as instruments, following Agan et al. (2023), who implement a similar approach by including a full set of judge dummies in the first stage. This specification serves as a check on our leave-one-out residualized JIVE (rJIVE) and

[37] Legal aid services were also available remotely and with limited in-person access.

CJIVE estimators, since judge-fixed effects use the same underlying variation but without the jackknife bias correction. Consistent with our main results, the estimated impact of granting a restraining order remains negative and statistically significant at the 10% level under conventional inference, and at the 1% level when standard errors are computed via a cluster bootstrap (894 daily filing-date clusters resampled with replacement).[38] We estimate a 17–19 percentage point reduction in the likelihood of a police-reported domestic violence crime at the address covered by the DVRO request, corresponding to an elasticity between -0.959 and -1.60.

In Columns (3)–(6), we implement the Influence-function Jackknife IV Estimator (IJIVE) and the Unbiased Jackknife IV Estimator (UJIVE). IJIVE, suggested by Chyn et al. (2025) and Agan et al. (2023), corrects for finite-sample bias that arises when a judge's own cases contribute to the construction of their leniency score, using a combination of leave-one-out and influence-function adjustments. UJIVE, introduced by Angrist and Kolesár (2024), further removes small-sample bias analytically, improving estimator coverage in settings with many instruments or limited clusters. Across both the IJIVE and UJIVE specifications, the estimated coefficients are similar to those from our main residualized JIVE specification, with magnitudes between –0.296 and –0.313. The weaker statistical significance (at the 10% level) suggests that IJIVE and UJIVE may not fully account for cluster-level biases (Frandsen et al., 2025).

Table 9 reports results from the saturated residualized JIVE specification, in which we fully interact the time fixed effects when residualizing the judge leniency measure. This approach allows us to verify that our results are not driven by the coarser calendar-date residualization used in the baseline model. The estimated effects remain highly consistent with the unsaturated JIVE model (-0.337 to -0.363). The pattern is stable across specifications using different temporal windows to construct the residualized JIVE instrument, suggesting that our main results are not sensitive to the particular time residualization scheme employed.[39]

In Table 10, we re-estimate the model without excluding judges who heard fewer than five cases in the previous year (peripheral judges). These low-volume judges are typically drawn from outside the dedicated family courts and may introduce additional variation in judicial behavior, as they are less likely to have received specialized training on restraining order procedures. Including them provides a robustness check that our results are not driven by the restriction to higher-volume judges. The CJIVE estimates remain qualitatively similar to the baseline, though somewhat larger in magnitude (a reduction of 0.37.9 to 0.46.2), suggesting that peripheral judges contribute additional identifying variation without materially changing the overall interpretation of the results.

Given the difference in the magnitude of the impact of granting a restraining order on the probability of a

[38] Agan et al. (2023) instead estimate their models using Limited Information Maximum Likelihood (LIML), taking advantage of their large sample size.

[39] Fully interacting the time fixed effects when constructing the CJIVE estimator would absorb nearly all of the day-level variation in judge assignment, leaving too little usable variation to identify the instrument. As a result, the CJIVE model cannot be reliably estimated under fully saturated time effects.

residential domestic violence crime or assault across restraining order types, one concern is that the lack of significance for civil restraining orders or DVROs without minor children might be driven by limited statistical power. To assess whether differences in power could explain the heterogeneous results, we computed minimum detectable effects (MDEs) for each subsample using a back-of-the-envelope calculation adapted for our judge-IV design.

Duflo et al. (2007) and Djimeu and Houndolo (2016) define the MDE as the smallest true effect that can be detected at a given significance level and statistical power, given the outcome variance and the variance of the exogenous treatment assignment. Their formulation uses t-statistics with finite-sample degrees of freedom, which is common in experimental design or small-sample settings where the intended sample size (N) and degrees of freedom are known:

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

where α denotes the rate of type I errors, β denotes the rate of type II errors, and P is portion of subjects allocated to treatment group. The statistical power is defined as $1 - \beta$.

In contrast to this experimental design framework, our observational setting assumes asymptotic normality with a large sample size, allowing the use of z-statistics. We also replace the square-root term with the corresponding regression standard error,[40] and replace the randomized treatment with our judge-leniency instrument.

In other words, this calculation is equivalent to multiplying the cluster-robust reduced-form standard error (SE_{ITT}) by $(z_{\text{power}} + z_{1-\alpha/2}) = 2.80$ for 80% power and $\alpha = 0.05$,[41] and then dividing by the absoluate value of the estimated first-stage coefficient π.

The reduced-form MDE is therefore

$$MDE_{ITT} = (z_{\text{power}} + z_{1-\alpha/2}) \times SE_{ITT},$$

The corresponding local average treatment effect (LATE) MDE divides the reduced-form MDE by the absolute value of the first-stage coefficient, π:

$$MDE_{IV} = \frac{MDE_{ITT}}{|\pi|}.$$

Table 11 reports the resulting minimum detectable effects across order types. Across subsamples, the implied MDEs are similar: the reduced-form MDEs range from 0.20 to 0.32 and the IV MDEs from 0.32 to 0.45 (80% power, 5% size, cluster-robust at the filing-date level). Because detectable effect sizes are comparable

[40] $SE_{ITT} = \sqrt{\frac{1}{P(1-P)} \cdot \frac{\hat{\sigma}^2}{N}}$ for a simple ITT under random assignment, homoskedastic outcomes, and no controls or clustering. However, because our analysis includes additional controls and clustered standard errors, this relationship does not hold exactly.

[41] For example, with 20 degrees of freedom, $t_{0.80} + t_{0.975} = 2.94$, which is slightly larger than the large-sample normal approximation.

across order types, differences in statistical power cannot explain why only DVROs with minor children yield statistically significant effects. As discussed in the next section, the differential impact likely reflects that DVROs with minor children often function as a bundled intervention—typically accompanied by custody orders granting sole custody or the removal of the respondent from the home—which enhances their deterrent impact.

5.3 Policy Relevance

Our 2SLS estimates capture a Local Average Treatment Effect (LATE): the causal effect of granting a restraining order on the probability of a subsequent crime for the subset of cases whose treatment status is influenced by judicial assignment. Under certain conditions, this quantity can also be interpreted as a Policy-Relevant Treatment Effect (PRTE), namely when a policy intervention shifts treatment probabilities along the same margin of unobserved resistance as the instrument (Heckman and Vytlacil 2001). In our context, the relevant policy lever is an institutional change in granting practices. Such a policy is not merely theoretical. The Los Angeles Superior Court (LASC) regularly conducts training for judges on handling restraining order requests. If training or other reforms systematically increased the granting rate, the resulting change in treatment would plausibly operate along the same margin as our judge-leniency instrument. This does not mean such a shift in the grant rate would necessarily produce the average deterrent effect in the newly granted orders. Should the grant rate be increased, we might expect to draw more marginal order requests into treatment, including petitions with different underlying characteristics and, potentially, deterrence effects. For example, suppose judges began granting more orders, we might expect to find grants of cases with lower probabilities of violence receiving grants.

We explore this question by estimating Marginal Treatment Effects (MTEs), which trace out how the causal effect of granting a DVRO varies across individuals with different unobserved propensities to receive treatment. MTE estimation requires stronger assumptions than conventional 2SLS, including a correctly specified model of treatment selection and adequate support over the unobserved resistance distribution. Our MTE estimation leverages our defendant and case characteristics (race, representation, and procedural indicators). Following Agan et al. (2023), we implement a control function approach to estimate the MTE, as developed in Heckman and Vytlacil (2007a) and Brinch et al. (2017).[42]

The top left panel of Figure 4 displays the distribution of estimated treatment propensities for treated (blue) and untreated (outlined) cases for all restraining order types. The vertical dashed lines indicate the trimmed region of common support, which spans propensity scores from approximately 0.1 to 0.65. We can only trace out the MTE along this range of support, so the substantial overlap across treated and untreated observations within this region suggests that treatment effects are identified for a broad and policy-relevant

[42]Specifically, we estimate the MTE using Stata's `mtefe` command (Andresen, 2018). Our specification uses a logit to estimate the propensity score. We trim 1% of the support to avoid extrapolation and estimate the MTE via local IV with a cubic polynomial specification, evaluated on the sample with common support.

range of marginal petitioners.

The right panel plots the estimated marginal treatment effect (MTE), measured with a cubic polynomial specification, as a function of unobserved resistance to treatment. The black line shows the point estimate of the treatment effect, while the shaded region denotes the 95% confidence interval. Conceptually, the x-axis of the MTE is an index of resistance to treatment: a latent measure of how likely a case is to be granted, conditional on observables. Cases with low resistance (near zero) are those that almost any judge would grant, such as petitions with strong evidence or clear histories of threat. Cases with high resistance (near one) are those that only the most lenient judges would grant, such as petitions with weaker documentation or procedural barriers. The MTE curve reveals whether deterrence is uniform across this distribution. The MTE declines monotonically, suggesting that DVROs may have larger deterrent effects for petitioners whose cases are relatively easy to grant and smaller effects for marginal cases that only lenient judges approve.

However, the estimates are imprecise, and none of the effects are statistically significant. Taken together, this suggests limited evidence of treatment-effect heterogeneity: granting appears broadly protective across the distribution, though the marginal gains from expanding the granting rate may be smaller than those implied by the average LATE. The picture is similar when we estimate the model using only DVROs with minor children in the middle panel of Figure 4. Consistent with our discussion of compliers, we find very little evidence of heterogeneous treatment effects.

6 Interpretation

Using a sample of restraining order requests from LASC's stand-alone restraining order courts, we estimate the impact of granting a civil restraining order on the probability that a residential, police-classified domestic-violence crime occurs at the block listed in the petition over the year following the 21-day hearing. In preliminary analyses, we find that DVROs with minor children are positively associated with subsequent crime, consistent with judges granting orders in higher-risk cases (endogeneity). By contrast, when we follow blocks for a year after the hearing, cases with a granted order show a declining daily crime probability relative to denied cases.

We then implement a judge-leniency instrumental variables design. For police-classified domestic violence crimes, granting a DVRO with minor children reduces the subsequent block-level probability of crime by 33–41 percentage points among cases whose treatment status is influenced by judicial assignment (the LATE). Results are similar when using aggravated assault as the outcome. Although the first stages remain strong for other order types, we do not detect any deterrent effect. The design exhibits similar statistical power across order types, indicating that the absence of significant effects for civil restraining orders and DVROs without minor children is unlikely to reflect limited power.

One theory to explain this pattern is that victims of intimate partner violence or threats from non-intimate

partners can take steps to protect themselves beyond the protection offered by a restraining order.[43] Even if the order is not granted, the petitioner can move, avoid interactions with the respondent, etc. For petitions involving minor children, the problem is more acute; due to the shared children, often the petitioner legally has to have some interaction with the respondent. If they do not get the order, it is much harder to avoid contact. The respondent, for example, would not necessarily be evicted from the shared home, making it more difficult, particularly for low-income petitioners who may not have easy access to other housing. DVROs also often include spousal and child support orders providing the petitioner court-mandated income, which reduces the respondent's bargaining power in Aizer (2010)'s model. In cases without a domestic violence restraining order, any dissolution of the relationship and the attendant custody and support orders would have to take place in the general family courts, which are much slower than the stand-alone restraining order courts that operate under the 21-day rule for restraining order hearings. More specifically, there may still be spousal support or child support orders in family court, but these will potentially take longer than a year. Moreover, absent a DVRO, separated parents are generally encouraged by courts to share custody or ensure that both parents have a meaningful relationship with their minor children, and custody orders reflect this aim.[44] In sum, absent a DVRO, the petitioner is likely to be forced by a custody order to continue to have some contact with the petitioner. However, if a DVRO is granted, the odds of joint custody or visitation decline dramatically. While granting a DVRO does not automatically or always result in sole custody being awarded to the petitioner, it significantly reduces them, as courts prioritize the safety and well-being of the child and the petitioner.[45] In cases where the court allows visitation, it is often supervised and involves minimal contact between the petitioner and the respondent. In short, a DVRO not only triggers additional sanctions if the respondent violates the order but also allows the petitioner an array of self-protection options. This theory does not explain the decline in the likelihood of a crime observed in the days following the issuance of an order. The decline is, however, consistent with Aizer and Dal Bo (2009)'s model of time-inconsistent preferences. Aizer and Bo note that the high rate at which domestic violence victims return to their abusers is consistent with a model of time-inconsistent preferences. Victims initially value their safety more than continuing their relationship, but over time, the disutility of continued violence is discounted relative to the contemporaneous value of the relationship. A decline in the probability of crime over time is consistent with this theory, as the granting of a DVRO makes it very difficult for a victim to allow the abuser back into their life and raises the stakes for the abuser if they do return. In Aizer and Bo's model, the DVRO grant is a commitment device preventing the victim from allowing the abuser to return.

[43] See, for example, Aizer (2010)'s household bargaining theory model. Aizer finds that higher relative wages for women reduce the likelihood of domestic violence.

[44] In LASC, as in many jurisdictions, there is a presumption favoring joint legal custody, where both parents share decision-making responsibilities, and even if one parent is given sole physical custody, the other parent is typically granted visitation rights unless there are concerns about the child's safety or welfare.

[45] In many jurisdictions, including California, there is a legal presumption against awarding custody to a parent who has perpetrated domestic violence within the last five years.

Yet this separation of children from the respondent makes the general welfare effects of civil restraining orders much more difficult to determine. Courts face a challenging balancing test. Set the threshold for granting an order too high, and more petitioners will be the victims of violence by their domestic partners. Setting the threshold too low can effectively sever the relationship between respondents and their children. This study cannot resolve this tension, but our estimates suggest that the tradeoff inherent in civil restraining orders is real, and the court's task is difficult.

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. (2010). The gender wage gap and domestic violence. *American Economic Review 100*(4), 1847–1859.

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.

Andresen, M. E. (2018). Exploring marginal treatment effects: Flexible estimation using stata. *The Stata Journal 18*(1), 118–158.

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.

Brinch, C. N., M. Mogstad, and M. Wiswall (2017). Beyond late with a discrete instrument. *Journal of Political Economy 125*(4), 985–1039.

Buzawa, E. S. and C. G. Buzawa (1996). *Do arrests and restraining orders work?* Sage Publications.

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.

Chaudhuri, M. and K. Daly (1992). Do restraining orders help? battered women's experience with male violence and legal process. In E. Buzawa and C. Buzawa (Eds.), *Domestic Violence: The Changing Criminal Justice Response*, pp. 227–231. Westport, CT: Auburn House.

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.

Heckman, J. J. and E. J. Vytlacil (2007a). Econometric evaluation of social programs, part i: Causal models, structural models and econometric policy evaluation. *Handbook of econometrics 6*, 4779–4874.

Heckman, J. J. and E. J. Vytlacil (2007b). Econometric evaluation of social programs, part ii: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. *Handbook of econometrics 6*, 4875–5143.

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.

Truman, J. L. and R. E. Morgan (2014). Nonfatal domestic violence. *Washington, DC: US Department of Justice, Bureau of Justice Statistics*.

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 Hearing

	(1) OLS	(2) OLS	(3) C-OLS	(4) rJIVE	(5) rJIVE	(6) rJIVE	(7) CJIVE	(8) CJIVE	(9) CJIVE
Panel A. Police-defined domestic violence									
RO Grant	-0.0131*	-0.0130+	-0.00624	-0.0744	-0.0669	-0.0535	0.00197	0.0111	-0.00354
	(0.00630)	(0.00661)	(0.00768)	(0.0725)	(0.0737)	(0.0771)	(0.0818)	(0.0842)	(0.0884)
AR 95% CI				[-.22,.068]	[-.215,.078]	[-.208,.098]	[-.184,.148]	[-.17,.171]	[-.178,.18]
Average Elasticity	-0.079	-0.078		-0.449	-0.404	-0.323	0.012	0.067	-0.021
Panel B. All Assault Continuum Incidents									
RO Grant	-0.00740	-0.00571	0.000247	-0.0888	-0.0853	-0.0760	-0.0162	-0.0110	-0.0422
	(0.00738)	(0.00779)	(0.00913)	(0.0826)	(0.0838)	(0.0871)	(0.0971)	(0.0999)	(0.104)
AR 95% CI				[-.255,.073]	[-.254,.079]	[-.252,.095]	[-.216,.175]	[-.217,.186]	[-.265,.157]
Average Elasticity	-0.033	-0.025		-0.395	-0.380	-0.338	-0.072	-0.049	-0.188
Panel C. Assault Continuum - Aggravated									
RO Grant	0.000807	0.00144	0.000647	-0.0343	-0.0352	-0.0405	0.00261	0.00188	-0.0118
	(0.00407)	(0.00432)	(0.00510)	(0.0463)	(0.0469)	(0.0507)	(0.0535)	(0.0551)	(0.0587)
AR 95% CI				[-.126,.058]	[-.129,.058]	[-.142,.06]	[-.108,.107]	[-.118,.106]	[-.128,.11]
Average Elasticity	0.014	0.024		-0.579	-0.593	-0.683	0.044	0.032	-0.199
Panel D. Assault Continuum - Simple									
RO Grant	-0.00821	-0.00715	-0.000401	-0.0545	-0.0501	-0.0355	-0.0188	-0.0129	-0.0304
	(0.00648)	(0.00692)	(0.00820)	(0.0722)	(0.0730)	(0.0750)	(0.0854)	(0.0875)	(0.0916)
AR 95% CI				[-.2,.087]	[-.197,.093]	[-.186,.112]	[-.198,.146]	[-.197,.157]	[-.23,.143]
Average Elasticity	-0.050	-0.043		-0.330	-0.303	-0.215	-0.113	-0.078	-0.184
Observations	20130	20079	18082	20129	20078	20080	20131	20080	20080
Control Variables	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
JIVE Years				Current&Last	Current&Last	Current&All Prior	Current&Last	Current&Last	Current&All Prior
First stage F-stat				146.587	144.845	134.046	74.299	64.59	57.762
First stage Coefficient				.69	.665	.65	.502	.469	.408
Complier Weights	No	No	Yes						

Notes: All specifications control for fixed effects. This table reports the main regression results by specifications, excluding cases dismissed before the restraining order hearing. Columns (1) and (2) report OLS specifications. Column (3) shows complier-weighted OLS results, splitting the sample into 10 mutually exclusive subgroups based on deciles of the predicted probability of granting order, estimated with all available covariates, following Agan et al. (2023). Columns (4) and (5) present coefficients from instrumental variable estimations using residualized JIVE with instruments from the current and previous year. Column (6) uses residualized JIVE with instruments from the current and all prior years. Columns (7) and (8) report CJIVE estimates based on the current and previous year, while Column (9) uses CJIVE with instruments from the current and all prior years. Each panel represents a distinct set of crime data. Panel A focuses on residential crimes that are associated with crime types found in the police-defined domestic violence incident dataset. Panel B includes the continuum of criminal assaults; Panel C focuses on aggravated assaults within Panel B; and Panel D examines simple assaults, threats, and child abuse from Panel B, excluding 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, + 0.10 p<0.1.

Table 5: Results by Restraining Order Type

	(1) OLS	(2) OLS	(3) rJIVE	(4) rJIVE	(5) rJIVE	(6) CJIVE	(7) CJIVE	(8) CJIVE
Panel A: Civil Harassment Order								
RO Grant	-0.00428	-0.00705	0.0212	0.0358	0.0530	0.0801	0.0895	0.0455
	(0.0103)	(0.0110)	(0.138)	(0.136)	(0.138)	(0.161)	(0.167)	(0.165)
AR 95% CI			[-.265,.299]	[-.247,.307]	[-.232,.333]	[-.266,.407]	[-.285,.422]	[-.337,.363]
Average Elasticity	-0.028	-0.046	0.137	0.231	0.343	0.518	0.579	0.294
Observations	9417	9406	9416	9405	9406	9417	9406	9406
First stage F-stat			44.985	49.864	48.861	31.694	27.447	27.328
First stage Coefficient			.531	.511	.515	.383	.335	.365
Panel B: DVRO								
RO Grant	-0.0194*	-0.0167+	-0.141	-0.152	-0.139	-0.0925	-0.100	-0.0849
	(0.00853)	(0.00906)	(0.0907)	(0.0928)	(0.0980)	(0.104)	(0.107)	(0.113)
AR 95% CI			[-.325,.038]	[-.34,.031]	[-.338,.055]	[-.308,.121]	[-.321,.12]	[-.318,.151]
Average Elasticity	-0.107	-0.092	-0.777	-0.835	-0.767	-0.509	-0.552	-0.467
Observations	9305	9297	9305	9297	9298	9306	9298	9298
First stage F-stat			97.335	98.029	86.324	40.751	40.998	35.467
First stage Coefficient			.81	.787	.755	.588	.633	.584
Panel C: DVRO with minor children								
RO Grant	0.00375	0.00386	-0.340*	-0.337*	-0.363*	-0.391*	-0.379*	-0.415+
	(0.0151)	(0.0158)	(0.167)	(0.166)	(0.184)	(0.197)	(0.191)	(0.216)
AR 95% CI			[-.743,-.039]	[-.734,-.038]	[-.824,-.037]	[-.919,-.047]	[-.871,-.042]	[-1.014,-.042]
Average Elasticity	0.020	0.021	-1.859	-1.843	-1.987	-2.139	-2.076	-2.270
Observations	3146	3142	3146	3142	3142	3146	3142	3142
First stage F-stat			33.992	36.108	28.913	21.056	24.869	19.549
First stage Coefficient			.779	.776	.714	.562	.606	.541
Panel D DVRO no children								
RO Grant	-0.0319**	-0.0282*	-0.0250	-0.0341	-0.00926	0.0895	0.0930	0.119
	(0.0107)	(0.0116)	(0.107)	(0.112)	(0.117)	(0.129)	(0.140)	(0.145)
AR 95% CI			[-.231,.203]	[-.25,.205]	[-.219,.259]	[-.156,.393]	[-.172,.431]	[-.15,.485]
Average Elasticity	-0.176	-0.156	-0.138	-0.188	-0.051	0.493	0.513	0.656
Observations	6159	6155	6159	6155	6156	6160	6156	6156
First stage F-stat			63.76	63.452	58.418	27.2	23.882	21.75
First stage Coefficient			.825	.791	.774	.602	.642	.599
Control Variables	No	Yes	No	Yes	Yes	No	Yes	Yes
JIVE Years			Current&Last	Current&Last	Current&All Prior	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	(6) CJIVE	(7) CJIVE	(8) CJIVE
Panel A: Police-defined domestic violence								
RO Grant	0.00375	0.00386	-0.340*	-0.337*	-0.363*	-0.391*	-0.379*	-0.415+
	(0.0151)	(0.0158)	(0.167)	(0.166)	(0.184)	(0.197)	(0.191)	(0.216)
AR 95% CI			[-.743,-.039]	[-.734,-.038]	[-.824,-.037]	[-.919,-.047]	[-.871,-.042]	[-1.014,-.042]
Average Elasticity	0.020	0.021	-1.859	-1.843	-1.987	-2.139	-2.076	-2.270
Panel B: All Assault Continuum Incidents								
RO Grant	0.00579	0.0104	-0.279	-0.285	-0.287	-0.303	-0.296	-0.324
	(0.0166)	(0.0174)	(0.176)	(0.178)	(0.196)	(0.204)	(0.201)	(0.227)
AR 95% CI			[-.687,.051]	[-.7,.043]	[-.758,.075]	[-.826,.072]	[-.797,.071]	[-.925,.087]
Average Elasticity	0.024	0.043	-1.154	-1.183	-1.188	-1.255	-1.228	-1.342
Panel C: Assault Continuum - Aggravated								
RO Grant	-0.00120	0.000205	-0.121	-0.124	-0.148	-0.199+	-0.195+	-0.208+
	(0.00938)	(0.00985)	(0.0860)	(0.0853)	(0.0971)	(0.103)	(0.0996)	(0.116)
AR 95% CI			[-.315,.044]	[-.317,.038]	[-.376,.036]	[-.465,-.011]	[-.441,-.012]	[-.509,.004]
Average Elasticity	-0.020	0.003	-2.036	-2.084	-2.488	-3.348	-3.275	-3.502
Panel C: Assault Continuum - Simple								
RO Grant	0.00699	0.0102	-0.158	-0.162	-0.139	-0.104	-0.102	-0.116
	(0.0152)	(0.0161)	(0.152)	(0.153)	(0.167)	(0.172)	(0.169)	(0.189)
AR 95% CI			[-.498,.137]	[-.504,.131]	[-.52,.184]	[-.511,.239]	[-.494,.23]	[-.575,.254]
Average Elasticity	0.038	0.056	-0.866	-0.889	-0.763	-0.570	-0.559	-0.636
Observations	3146	3142	3146	3142	3142	3146	3142	3142
Control Variables	No	Yes	No	Yes	Yes	No	Yes	Yes
JIVE Years			Current&Last	Current&Last	Current&All Prior	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 with a minor child identified as present and excluding civil harassment prevention cases and domestic violence restraining order petitions without a minor child present. We start Panel C of Table 4. 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 7: Excluding Virtual Period (Domestic Violence Restraining Order with Minor Child)

	(1) OLS	(2) rJIVE	(3) rJIVE	(4) CJIVE	(5) CJIVE
RO Grant	0.00399	-0.333^{+}	-0.355^{+}	-0.377^{+}	-0.412^{+}
	(0.0161)	(0.179)	(0.196)	(0.212)	(0.238)
AR 95% CI		[-.764,-.011]	[-.852,-.007]	[-.935,-.002]	[-1.095,-.001]
Observations	3045	3045	3045	3045	3045
Average Elasticity	0.022	-1.813	-1.936	-2.051	-2.245
Control Variables	Yes	Yes	Yes	Yes	Yes
JIVE Years		Current&Last	Current&All Prior	Current&Last	Current&All Prior

Notes: This table replicates the regression results from Table 6, focusing on Panel A. It excludes petitions filed from March 15, 2020, when the restraining order courts went virtual, to June 5, 2020, when the standalone restraining order courts returned to in-person hearings. It retains the full set of judges rather than excluding peripheral judges, defined as hearing fewer than five cases in the previous year, for statistical power. 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 8: Alternative Instrument Measures (Domestic Violence Restraining Order with Minor Child)

	(1)	(2)	(3)	(4)	(5)	(6)
	Dummy	Dummy	IJIVE	IJIVE	UJIVE	UJIVE
RO Grant	-0.175**	-0.194**	-0.296+	-0.285	-0.301+	-0.313+
	(0.062)	(0.066)	(0.161)	(0.183)	(0.170)	(0.184)
AR 95% CI	[-.396,.026]	[-.439,.024]	[-.682,-.004]	[-.739,.048]	[-.705,.008]	[-.768,.019]
Observations	3142	3142	3142	3142	3142	3142
Average Elasticity	-0.959	-1.060	-1.621	-1.557	-1.646	-1.714
Control Variables	Yes	Yes	Yes	Yes	Yes	Yes
JIVE Years	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. 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. For Columns 1 and 2, standard errors are computed via a cluster bootstrap (894 daily filing-date clusters resampled with replacement) to account for the fitted-instrument setting with many judge dummies as instruments.
*** p<0.001, ** p<0.01, * p<0.05, + 0.10 p<0.1.

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: CJIVE Preserving All Judges (Domestic Violence Restraining Order with Minor Child)

	(1) CJIVE	(2) CJIVE	(3) CJIVE	(4) CJIVE
RO Grant	-0.379*	-0.415^{+}	-0.462*	-0.390^{+}
	(0.191)	(0.216)	(0.209)	(0.236)
AR 95% CI	[-.871,-.042]	[-1.014,-.042]	[-1.056,-.105]	[-1.077,.023]
Average Elasticity	-2.076	-2.270	-2.528	-2.132
Observations	3142	3142	3135	3135
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 retains the full set of judges rather than excluding peripheral judges, defined as hearing less than five petitions in the last year. 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 11: 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 $\text{MDE}_{ITT}/|\pi|$.

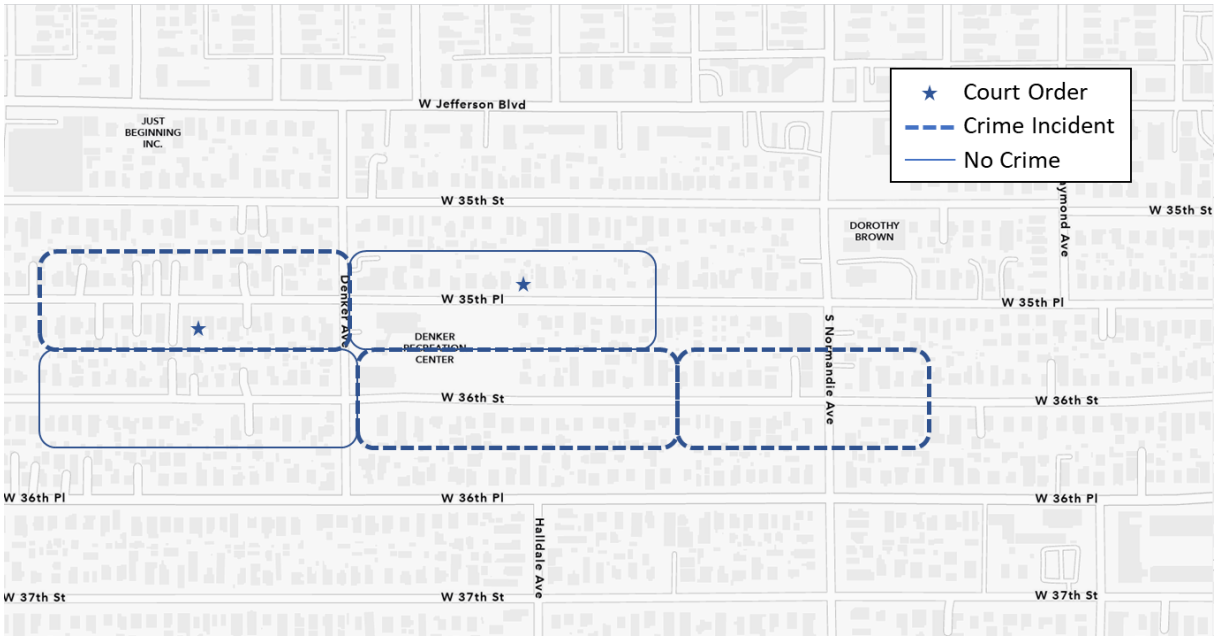
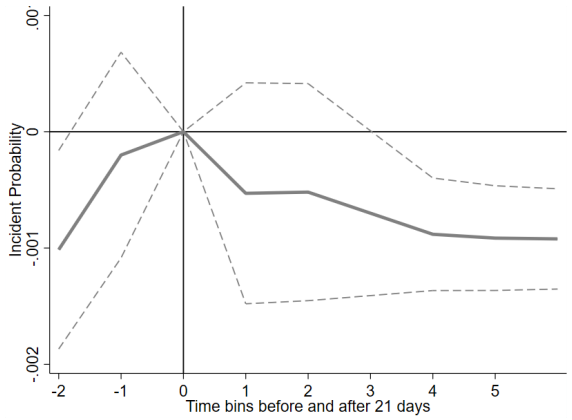
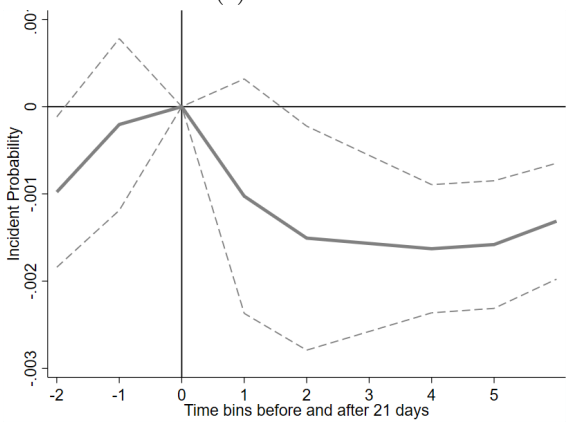


Figure 1: Hypothetical Examples of 100-block Address Match

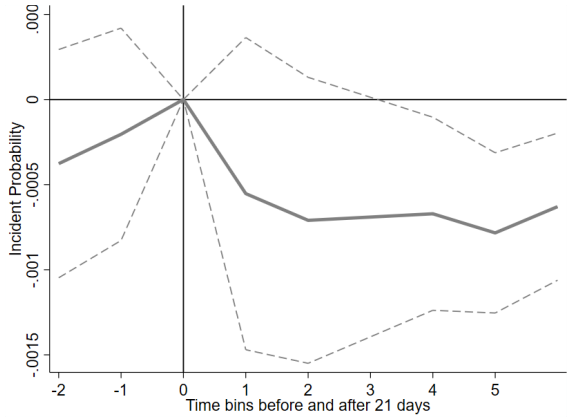
Note: This figure illustrates the matching process for addresses aggregated to the 100-block level. Black stars are hypothetical addresses similar to those found in restraining order requests, blue boxes show 100 blocks, and dashed lines indicate a hypothetical crime on that street.



(a) Police DV



(b) Assault Continuum



(c) Aggravated Assault

Figure 2: Impact of Restraining Orders on Daily Crime Probability: Event Study Analysis

Note: This figure illustrates the event study analysis of domestic violence restraining orders involving minor children, using weekly time bins relative to the 21-day hearing date as the event point (week 3), and ignoring petitions dismissed before the 21-day hearing. The analysis estimates the probability of a crime occurring on a given day within the 100-block of the address listed in the restraining order, with separate trends for granted and denied orders. The reference period is week 3, aligning with the restraining order hearing. The model controls for case and year-month fixed effects, with standard errors clustered by both order request and date.

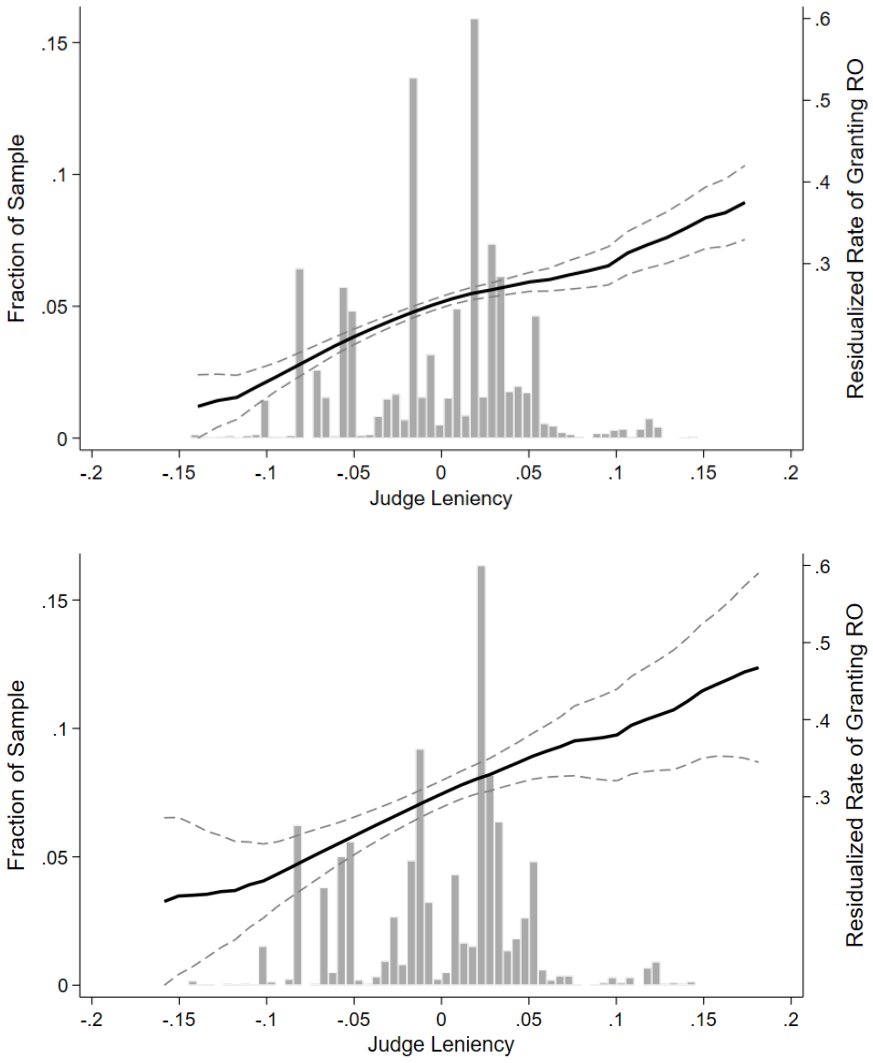


Figure 3: Distribution of Residualized Leniency Measure and First Stage

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.

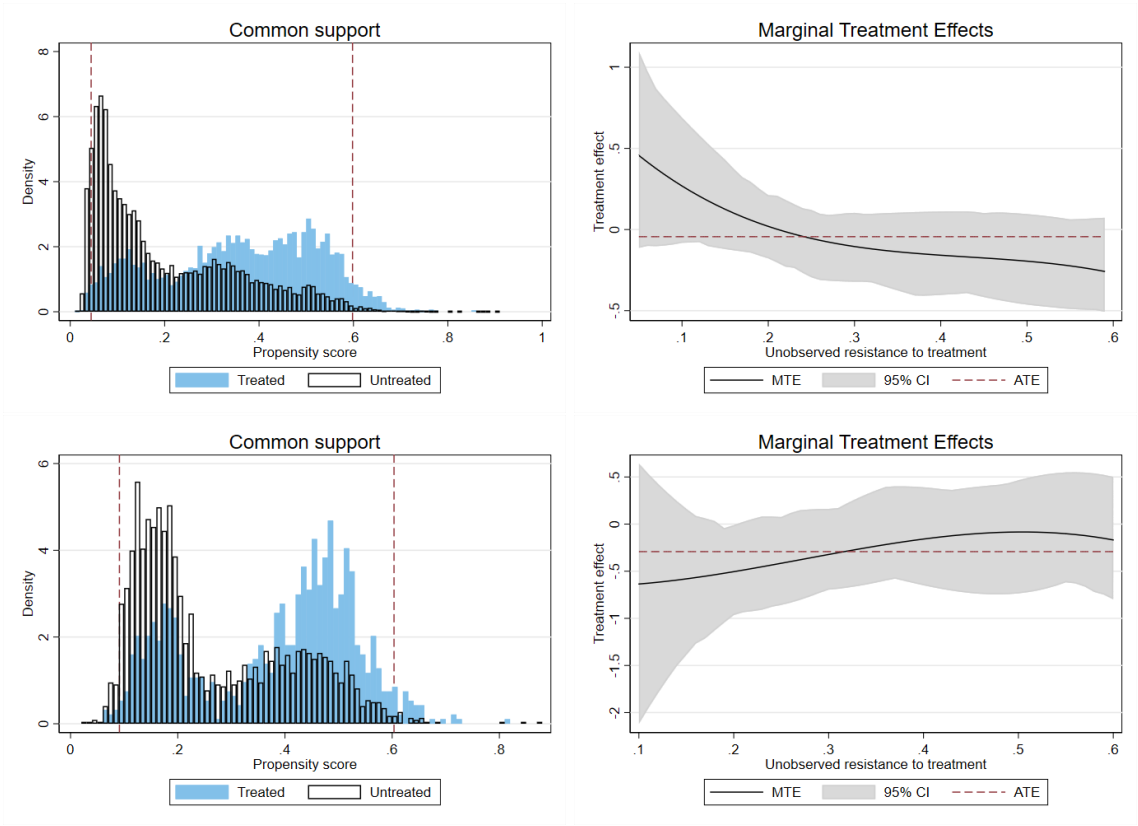


Figure 4: Common Support and MTE Plot

Note: Panel (a) plots the density of the estimated propensity score predicted by a logit regression, including the covariates listed in the text and fixed effects for year, month, day of week, and case. Vertical dashed lines indicate the 1% trimmed common-support bounds used for estimation. Panel (b) shows the marginal treatment effect (MTE) estimated via local IV with a cubic polynomial specification, evaluated on the sample with common support. The x-axis is the unobserved resistance to treatment $u \in [0, 1]$ (the rank of the latent index). In this context, resistance reflects unobserved case characteristics that make a DVRO more or less likely to be granted, even after conditioning on observables and fixed effects. Confidence intervals are based on 100 bootstrap replications; standard errors are clustered by filing date. The outcome is an indicator for a police-identified domestic violence incident at the block listed in the restraining order request. Estimates were obtained with `mtefe` in Stata (Andresen, 2018). The estimation uses only requests for DVROs with minor children.

A Additional Tables

Table A1: Event Study of Probability of a Crime

	(1) Police-defined DV	(2) Assault Continuum	(3) Aggravated Assault	(4) Simple Assault
Grant	0.000698***	0.00135***	0.000647**	0.000720**
	(0.000187)	(0.000298)	(0.000209)	(0.000231)
Week 1 × Grant	-0.00101*	-0.000979*	-0.000377	-0.000599^{+}
	(0.000427)	(0.000431)	(0.000335)	(0.000333)
Week 2 × Grant	-0.000201	-0.000205	-0.000204	0.00000222
	(0.000442)	(0.000492)	(0.000312)	(0.000334)
Week 4 × Grant	-0.000530	-0.00103	-0.000553	-0.000466
	(0.000475)	(0.000671)	(0.000458)	(0.000381)
Week 5 × Grant	-0.000520	-0.00151*	-0.000710^{+}	-0.000975*
	(0.000466)	(0.000642)	(0.000420)	(0.000453)
Week 6 × Grant	-0.000882***	-0.00163***	-0.000671*	-0.000662^{+}
	(0.000242)	(0.000367)	(0.000283)	(0.000363)
Week 7 × Grant	-0.000915***	-0.00158***	-0.000784**	-0.000782**
	(0.000225)	(0.000365)	(0.000235)	(0.000262)
Week 8 × Grant	-0.000922***	-0.00131***	-0.000629**	-0.000576*
	(0.000215)	(0.000332)	(0.000216)	(0.000238)
Observations	1,123,254	1,123,254	1,123,254	1,123,254
Year-month FE	Yes	Yes	Yes	Yes
Two-way cluster	Yes	Yes	Yes	Yes

Notes: This table reports event-study estimates of the effect of domestic violence restraining order (RO) grants on the daily probability of a domestic-violence–related crime at the block level. The event date is defined as the 21-day hearing at which the temporary restraining order is either granted or denied. Each coefficient measures the difference in the probability of a crime for addresses associated with granted versus denied petitions in the indicated week relative to the reference period (Week 3 × Grant). The model includes case fixed effects and year–month fixed effects. Standard errors are two-way clustered by order request and date.$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.