Justice as Deterrence: The Effect of Convictions on Domestic Violence in India*

Daniel L. Chen[1], Anders Kjelsrud[2], and Balasurya Sivakumar[2]

[1]Toulouse School of Economics

[2]Oslo Metropolitan University

March 10, 2026

Abstract

Can legal enforcement deter behavior that is private, underreported, and shaped by social norms? We study this question in the context of intimate partner violence in India, where abuse largely occurs behind closed doors and observed crime statistics conflate true incidence with reporting behavior. We combine nationally representative survey data that retrospectively measure the onset of domestic violence with administrative records on reported crimes and the universe of lower-court cases involving crimes against women from 2010–2018. To address endogeneity in enforcement, we exploit quasi-random assignment of criminal cases to judges to construct district–year variation in conviction probabilities driven by differences in judicial harshness. This variation shifts the expected likelihood of punishment while holding laws, policing resources, and court infrastructure fixed. Our instrumental-variables estimates show that higher conviction probabilities causally reduce subsequent violence: a one-percentage-point increase in conviction rates lowers the probability of violence onset in the following year by 0.22 percentage points, corresponding to 15 percent of the sample mean. At the same time, reported crimes against women increase following higher conviction rates. We show that this divergence reflects increased willingness to report rather than higher victimization, highlighting how enforcement can simultaneously deter violence and raise measured crime. Consistent with a salience-based mechanism, deterrent effects are significantly larger in areas with greater exposure to local media, particularly community radio and local newspapers. Together, the findings demonstrate that visible adjudication can deter private violence and empower victims to engage with formal institutions, while underscoring the importance of distinguishing behavioral change from reporting responses when evaluating the effects of law.

*

1 Introduction

A central challenge in the economics of law is determining whether legal enforcement can deter behavior that is private, underreported, and shaped by social norms. In such settings, standard measures of crime confound changes in true behavior with changes in reporting, and observed enforcement outcomes are themselves endogenous to underlying social conditions. As a result, even when legal institutions expand or punishments increase, it remains unclear whether the law reduces harmful behavior or merely alters how that behavior is recorded.

Intimate partner violence represents an extreme case of this challenge. Abuse typically occurs behind closed doors, victims face substantial social and economic barriers to reporting, and formal enforcement often lacks visibility. In India, roughly one in three women report experiencing physical or sexual abuse—well above the global average (NFHS, 2021)—despite the existence of criminal statutes designed to protect women. The Protection of Women from Domestic Violence Act of 2005 expanded legal protections nationwide, yet enforcement and adjudication continue to vary widely across states, districts, and courts. This variation raises a fundamental question: can criminal convictions deter domestic violence in contexts where behavior is hidden and reporting is endogenous?

Theory offers competing predictions. On the one hand, higher conviction probabilities may deter abuse by raising the expected cost of violence and by signaling stronger social disapproval. On the other hand, more active enforcement may increase trust in formal institutions and encourage victims to report, potentially increasing observed crime even if true violence declines. Distinguishing between these mechanisms is empirically difficult, as enforcement affects both behavior and measurement.

We address this challenge by combining three sources of data that allow us to separately identify changes in actual violence and changes in reporting behavior. First, we use nationally representative survey data from the National Family Health Survey, which includes retrospective information on the timing of women's first experience of intimate partner violence. This enables us to construct a measure of violence onset rather than relying solely on contemporaneous prevalence. Second, we use administrative data from the National Crime Records Bureau on reported crimes against women. Third, we compile administrative records on the universe of lower-court cases involving crimes against women from 2010–2018, including filing dates, decision dates, and judge identifiers.

Our empirical strategy exploits quasi-random assignment of criminal cases to judges in India. We construct a district–year measure of expected conviction probabilities based on the composition of judges deciding cases, using out-of-sample variation in judicial harshness derived from judges' prior rulings. Because cases are assigned to judges as-good-as randomly within courts, this variation shifts the likelihood of conviction without changing statutes, policing resources, or local enforcement capacity. This allows us to isolate the causal effect of adjudication itself on subsequent violence.

We document three main findings. First, higher conviction probabilities causally reduce subsequent domestic violence: a one-percentage-point increase in conviction rates lowers the probability of violence onset in the following year by 0.22 percentage points, approximately 15 percent of the sample mean. Second, despite this reduction in survey-measured violence, reported crimes against women increase following higher conviction rates. We argue that this divergence reflects greater willingness to report rather than higher victimization, illustrating how enforcement can simultaneously deter harmful behavior and raise recorded crime. Third, deterrent effects are significantly larger in areas with greater exposure to local media, particularly community radios and local newspapers, consistent with visibility and information diffusion as key preconditions for deterrence.

This paper makes three contributions. First, we provide causal evidence that criminal adjudication can deter private, underreported violence, even in settings with weak enforcement capacity. Second, we show that increases in enforcement can reduce true violence while increasing reported crime, offering a framework to reconcile conflicting findings in the empirical literature. Third, we demonstrate that deterrence depends critically on visibility: identical increases in conviction risk have larger effects where legal outcomes are more salient, highlighting the role of information diffusion and belief updating in legal compliance.

We contribute to a growing literature on whether law can deter behavior that is private, underreported, and shaped by social norms. As mentioned, a central challenge in this literature is that observed crime statistics typically conflate changes in true behavior with changes in reporting, while enforcement outcomes are themselves endogenous to both. Much of the existing evidence focuses on legal reforms—such as domestic-violence statutes and divorce laws—that simultaneously alter legal rights, reporting channels, and exit options. These studies generally find reductions in violence, either by lowering the cost of leaving abusive relationships or by raising the expected cost of abuse (Beleche, 2019; Bobonis et al., 2013; Corradini and Buccione, 2023; Ferraz and Schiavon, 2022; García-Ramos, 2021; Sanin, 2025). In South Asia, however, evidence remains limited and mixed: Gulesci et al. (2024) find no average effect of a domestic-violence law in Pakistan and document increases in violence in more conservative areas.[1] We contribute by isolating the effect of criminal adjudication itself and by showing how enforcement can simultaneously reduce true violence while in-

[1]Beyond domestic-violence statutes, related work studies reforms to property and inheritance rights (Amaral, 2017; Anderson, 2021; Harari, 2019). These reforms can reduce violence through improved bargaining power but may also provoke backlash, underscoring the difficulty of interpreting observed changes in reported abuse.

creasing observed crime. Crucially, we argue that deterrence in such private and stigmatized settings depends on the visibility of legal enforcement: convictions affect behavior not only through formal sanctions, but by shaping beliefs about the likelihood and social meaning of punishment.

Our analysis also speaks to the literature on policing and enforcement capacity in low- and middle-income countries, where weak institutions complicate deterrence and measurement. A prominent policy response has been the creation of specialized enforcement units—such as women police stations or women justice centers—intended to lower reporting costs and improve responsiveness to crimes against women. The evidence from these interventions is mixed and highly context dependent. In India and Brazil, women police stations increased reporting but did not reduce domestic violence or homicides on average, though they reduced homicides in large Brazilian cities (Amaral et al., 2021; Perova and Reynolds, 2017). In contrast, women justice centers in Peru reduced violence while also increasing reporting (Sviatschi and Trako, 2024). These findings highlight a core identification problem: increases in recorded crime following enforcement improvements may reflect greater willingness to report rather than higher victimization. Our design directly addresses this issue by separately measuring changes in violence incidence and reporting behavior, allowing us to distinguish deterrence from measurement effects. These mixed findings suggest that enforcement capacity alone is insufficient: when legal action is not salient or widely understood, improvements in policing may primarily affect reporting rather than behavior. Our design allows us to test this directly by examining whether identical increases in conviction risk have larger deterrent effects in environments where legal outcomes are more visible.

Finally, the paper contributes to the economics-of-law literature that leverages the digitization of court records to study judicial decision-making and its real-world consequences (Ramos-Maqueda and Chen, 2025). A common strategy in this literature is to exploit quasi-random assignment of judges to cases to identify causal effects of judicial rulings. Prior work has used such designs to study incarceration, bail decisions, and downstream outcomes such as employment and recidivism (Aizer and Doyle, 2013; Arnold et al., 2018; Bhuller et al., 2020). In India, the expansion of the eCourts system and related digitization efforts has enabled large-scale analysis of judicial behavior.[2] Existing studies find little evidence of in-group bias among judges and document links between judicial rulings and environmental outcomes such as air and water pollution (Behrer et al., 2025; Bhupatiraju et al., 2024; Do et al., 2018). To the best of our knowledge, no prior study uses quasi-random judicial assignment to examine whether criminal adjudication deters domestic violence or to show

[2] *IndiaKanoon* is a search engine for Indian cases and statutes. Ash et al. (2025) construct a comprehensive dataset covering lower-court cases across India.

that the effectiveness of adjudication depends on the visibility of legal outcomes. By linking judicial decisions to local media exposure, we demonstrate that information diffusion is a central mechanism through which courts influence behavior outside the courtroom.

2 Background and the Domestic Violence Act of 2005

2.1 Domestic violence

India is among the most gender-unequal countries, ranking 108 out of 193 in the Global Gender Inequality Index (UNDP, 2022). Rigid social and cultural norms contribute to discrimination across multiple dimensions, leaving women behind men in adult literacy, employment, asset ownership, and household bargaining power. This inequality extends into households and intimate relationships, where 35% of women report having experienced intimate partner violence.[3] This is 8 percentage points higher than the global average of 27%. According to the NFHS 2021 report, 28% of women face less severe physical violence, 14% experience emotional or psychological violence, 9% suffer severe physical violence, and 6% are subjected to sexual violence.

Until 2005, domestic violence complaints were prosecuted under the anti-dowry and anti-cruelty provisions of the Indian Penal Code. Both statutes were considered inadequate, as they were applied only by the police in cases involving extreme cruelty, such as those leading to suicide, or in relation to dowry disputes. Women's rights groups across the country therefore demanded a new law to address non-dowry-related forms of domestic abuse (Vyas, 2006). This culminated in the passage of the Protection of Women from Domestic Violence Act (PWDVA) in September 2005. The Act provides comprehensive protection against physical, emotional, financial, psychological, and sexual violence in both marital and non-marital relationships. It also obliges the state governments to appoint protection officers, provide free legal aid, and ensure access to counselors and medical facilities to support victims of domestic violence.

The implementation of this law has been largely uneven across Indian states. UN Women (2014) highlights that only one-third of the states have made any budgetary allocation under the PWDVA to provide for dedicated protection officers, medical facilities and legal aid for victims of domestic violence. These allocation also vary a lot between states ranging from USD 3000 in Meghalaya to USD 800,000 in Karnataka (Das and Lakshmana, 2020). Similarly, the appointment of protection officers differs widely across states ranging from 3760 officers in Maharashtra to only 33 in Tamil Nadu. Furthermore, the Supreme Court

[3] See here: https://genderdata.worldbank.org/en/economies/india

of India has repeatedly reprimanded states for not implementing various provisions of the PWDVA.[4]

2.2 Court system in India

India follows a common-law system with a hierarchical court structure. At the apex is the Supreme Court, followed by state High Courts, and then the district and subordinate courts that together constitute the lower judiciary. We focus on these lower courts because they are the point of entry for almost all criminal cases, with appeals moving upward only after an initial judgment. India follows a common-law tradition in which judicial decisions shape subsequent rulings through precedent, making judges particularly influential in the development of legal practice and enforcement (Ash et al., 2025). Moreover, criminal trials are decided by judges rather than juries, so the trial judge plays a central role in determining case outcomes.

2.3 Assignment of cases to judges

A criminal case reaches the courts after a First Information Report (FIR) is filed at a local police station. Because each police station is mapped to a particular district courthouse, territorial jurisdiction largely determines which courthouse receives a filing. The subsequent allocation of cases to judges within that courthouse is governed by administrative rules rather than litigant choice. In the simplest setting—when only one criminal judge is sitting—all cases are heard by that judge. In courthouses with multiple judges, filings are routed using pre-set rosters that link *both* the police station of origin and the statutory charge to a particular courtroom. The judges assigned to that courtroom at the relevant time then become responsible for the case.

Several institutional features further limit scope for selecting a preferred judge. Judges do not remain permanently tied to a single courtroom: over the course of their tenure at a courthouse, they rotate across different courtrooms and case portfolios, so the same police-station–charge category will be handled by different judges at different points in time. Moreover, the first hearing date and the progression of a case are uncertain at the time of filing due to congestion and delays, making it difficult to predict which judge will ultimately preside when key decisions are made. The judiciary also takes an explicit stance against attempts to steer cases toward particular judges ("judge shopping"), which further discourages manipulation (Ash et al., 2025; Mehmood et al., 2023).

[4] Supreme Court Observer, June 2025.

These institutional arrangements motivate our identifying assumption of *quasi-random* judge assignment conditional on court, time, and legal section, following the approach in Ash et al. (2025). Conditioning on the courthouse and time period captures the local pool of judges and administrative conditions, while conditioning on the legal section captures the statutory category that governs routing. Within these cells, residual variation in which judge hears a case is plausibly orthogonal to unobserved determinants of conviction.[5] In Subsection 4.3, we assess whether the assignment of judges to cases in our data appears consistent with quasi-random allocation. In line with the evidence documented in Ash et al. (2025), we find no indication that observable case characteristics predict the assigned judge, supporting our identification strategy.

3 Data

Our data comprise three main components: (i) measures of violence and crimes against women, (ii) detailed information on court cases, and (iii) geographical data on local media outreach.

Violence and crimes and against women. We use two different data sources to capture violence and crimes against women.

Our first source is the 2015–16 (NFHS-4) and 2019–21 (NFHS-5) waves of the Indian National Family Health Survey. Each wave interviewed approximately 700,000 women aged 15–49. A subsample of 10% of these women also completed a domestic-violence module, which covers experiences of physical and sexual violence. Among women who completed this module, about 31% reported having experienced violence.[6] Women who report violence are additionally asked when the first episode occurred, measured in terms of years since marriage. Because the survey records the date of marriage, we are able to use this to construct a person-year panel that traces each woman from marriage until the year in which she first experiences violence. In Figure 1, we present a hazard plot of first violence since marriage. For each marriage-duration year, the plot shows the share of women who experience a first episode of

[5] Additionally, in the Indian context, only about 0.05% of criminal cases result in a plea bargain, thereby alleviating concerns about selection bias arising from the types of cases that proceed to trial (Ash et al., 2025).

[6] Females are considered to have experienced violence if they affirm to either one of the following violent scenarios: i) *pushing*, ii) *twisting of an arm or pulling hair*, iii) *slapping*, iv) *punch with a fist or something that could hurt*, v) *kicking or dragging*, vi) *choking or burning*, vii) *attacking with a knife, gun or other weapon*, viii) *physically forced sexual intercourse*, ix) *physically forced other sexual acts* and x) *forced sexual acts through other threats*.

violence in that year among those with no prior episode.[7] The figure indicates that violence typically begins within the first few years of marriage and rarely begins more than ten years after marriage. Because of this, we also estimate a specification on a sample limited to women no more than ten or five years into marriage at the time of the survey.[8]

Our second data source is formally reported from the National Crime Records Bureau (NCRB), a government agency under the Ministry of Home Affairs in India. We use district-wise data on crimes committed against women, the most disaggregated level available, from the NCRB's annual *Crime in India* reports for the years 2011–2022. The NCRB reports crimes under both the Indian Penal Code (IPC) and Special and Local Laws (SLL). The IPC covers violent crimes such as rape, dowry deaths, and acid attacks, while SLL include offenses under special legislations such as the Protection of Women from Domestic Violence Act (PWDVA), 2005, and the Immoral Traffic (Prevention) Act, 1956, aimed at countering human trafficking. For our analysis, we rely on aggregated measures of crimes against women, as crime statistics under several special laws are not consistently reported across all years. Finally, it is important to note that the NCRB data capture only reported crimes and therefore may suffer from under-reporting, a well-documented concern in the context of gender-based violence in India.

Court cases. We use case-level records from India's eCourts platform covering filings from 2010–2018, as compiled by Ash et al. (2025). The India eCourts platforms was established by the Indian Government in 2013 to digitize court records as part of a broader initiative to modernize the judiciary with IT capabilities. These digitized court records contained full text orders and judgments of courts cases along with other useful information. Ash et al. (2025) compiles this data by obtaining 77 million cases covering courts in subordinate judiciary — courts at and below the District and Sessions level where all the criminal cases are first filed. This compilation by Ash et al. (2025) reports key timestamps for each case (filing, hearing, and decision dates and unique judge identifiers), as well as information on whether the crime under trial was committed against a women.

Using this information we only include cases on crimes against women filed under the Indian Penal Code or the Code of Criminal Procedure as we are interested in exploring the effect of convictions in crimes against women.

[7]Note that the nature of the dataset implies that women exit the panel in the year of their first episode of violence.

[8]Another reason for restricting the sample is that women tend to heap their reporting for violence onsets many years after marriage (10, 15 and 20).

Local media. We collect data on local media to explore the heterogeneity in treatment effects based on media penetration across India. We compile three types of measures.

First, we use data on self-reported media use from the 2015–16 wave of the NFHS. Second, we obtain information on community radio stations from a circular issued by the Broadcasting Wing of the Ministry of Information and Broadcasting (Minster of Broadcasting and Information, 2020). This circular has detailed information on the address and the opening date of all the 289 community radio stations in India that are operating in India as of 2020[9]. We geocode each stations address using the Google Geocoding API to obtain latitude and longitude coordinates.[10] Figure 2 displays the spatial distribution of community radios across India.

Third, we complement this with district-level data from Registrar of Newspapers in India (RNI) for the year 2013-14 as digitized and compiled by Cage et al. (2023) from the *Press in India* reports. The RNI is a public authority that is tasked with registering new newspapers and tracking circulation figures. In their annual Press in India reports, they provide information on the number of newspaper titles and total circulation across Indian cities, covering all languages and periodicity. Cage et al. (2023) allocate newspapers to districts based on their location, arguing that papers printed in the cities are circulated to the surrounding rural areas. We use their district-level data on the number of titles and circulation figures for the financial year 2013-2014, which is the start of our period of study.[11]

4 Empirical Approach: Judge IV

Conviction rates are unlikely to be exogenous: districts with higher convictions may differ systematically in unobserved ways that also affect domestic violence (e.g. policing intensity, prosecutorial effort, local social norms, or reporting behavior), and conviction rates may themselves respond to local crime trends. To address this endogeneity, we build an instrumental that isolates plausibly exogenous variation in convictions stemming from quasi-random assignment of cases to judges with different propensities to convict, as described below. The premise of the analysis is that cases are randomly assigned to judges conditional on court, time and charge (see Section 2).

[9] Our analysis uses stations operating as of 2020. Roughly 95% of stations observed in 2014 are also listed in 2020, suggesting limited exit.

[10] For 287 out of 289 radios we found that the coordinates generated lie in the same district as stated in the address. For the two remaining, one station could be manually located with information in the address and one station was excluded as a reliable location could not be identified with the available information.

[11] We are grateful for Guilhem Cassan and co-authors for sharing this data with us.

4.1 Building the instrument

To keep timing clear, let y denote the *filing* year (when the random assignment occurs), t the *decision* year (when the court decision materialize), and τ a rolling cutoff used to construct judge harshness from *pre-τ* cases.

Step 1: Judge-level harshness. For each cutoff $\tau \in \{2013, \ldots, 2018\}$, we restrict the sample to criminal cases against women filed in years $y < \tau$ and to judges with at least five such pre-τ cases. We then estimate

$$\text{Conviction}_{ijy} \;=\; \delta_{cy} \;+\; \delta_s \;+\; \varepsilon_{ijy}, \qquad y < \tau, \tag{1}$$

where Conviction_{ijy} equals one if case i handled by judge j ends in conviction, δ_{cy} denotes court-year fixed effects, while δ_s denotes charge section fixed effects. The charge section fixed effects ensure comparisons within similar charges, and the court–year fixed effects ensure comparisons within the same court–year (this is the same set of fixed effects as used by Ash et al., 2025). Using this, we next define judge j's expected harshness in year τ as the mean residual

$$H_{j\tau} \;=\; \frac{1}{N_{j,<\tau}} \sum_{y<\tau} \sum_{i} \varepsilon_{ijy}. \tag{2}$$

$H_{j\tau}$ is thus an *ex ante* measure that captures whether judge j in the past has tended to convict more or less than other judges, after accounting for common type of charge and court-year effects.

We could have included more granular fixed effects to align cases even more finely by time and charge. We do not further saturate the baseline, however, because the resulting cells often have little (or no) within-cell variation, sharply reducing the effective sample used to identify judge harshness. In Appendix B we still show that our main results are essentially unchanged — albeit somewhat less precisely estimated — when we (i) interact the court–year fixed effects with month, (ii) interact the court–year fixed effects with charge section, or (iii) interact the court–year fixed effects with both month and charge section (Table A4). For efficiency, we retain the specification in Equation 1, as any potential bias from not saturating further appears minimal in practice.

Step 2: District-level harshness. Our outcome variables can only be linked to the court data at the level of district-year. Hence, we cannot use the above measures directly in our analysis. Instead we want a measure that captures the mix of judges (harsh or lenient) a district ended up with in a particular year. Also, because we are interested in the potential preventive effect of convictions on violence, the relevant timing is the date of court *decisions*, not case filings.

We therefore proceed as follows. We first attach the harshness measure H_{jy} to court cases based on judge and filing year. For instance, if a case was filed in 2015 and decided in 2016 we attach $H_{j,2015}$, which captures the judge's harshness computed using court cases filed *prior* to 2015. We next aggregate based on districts and decision years

$$Z_{dt} \;=\; \sum_{j \in J_{dt}} w_{jdyt}\, H_{jy}, \tag{3}$$

where w_{jdyt} denote the share of cases decided in district d in year t that was filed to judge j in year y.

There are two potential issues with this way of aggregating. First, the randomization of cases happens at the time of filing, and only 49% of the cases in our data have identical filing and decision years. As a robustness check, we show that our results continue to hold when we restrict the analysis to this subsample (see Table A6), suggesting that the timing mismatch is not driving our findings. Second, judges might move between courts and the variation in our instrument could therefore be dominated by which judges are available rather than by the (random) allocation of cases. As another test, we simulate a counterfactual random assignment of cases to the set of judges available at each point in time and re-estimate the judge-harshness measure. This placebo measure has no predictive power for conviction rates, indicating that the identifying variation in our main instrument arises from actual case allocation, not judge availability (see Figure A1).

4.2 Main specification

Our main specification takes the following form

$$Y_{dt} \;=\; \beta\, \widehat{\text{ConvictRate}}_{d,t-1} \;+\; \delta_d \;+\; \lambda_t \;+\; \varepsilon_{dt}, \tag{4}$$

where Y_{dt} denotes either the survey based violence measure, or formally reported crimes against women in year t. $\widehat{\text{ConvictRate}}_{d,t-1}$ is the predicted conviction rate of crimes again women in district d in year $t-1$. δ_d and λ_t denote district and year fixed effects, respectively. For the survey based violence measure, for which we have data at the individual level, we also add a set of marriage year (length) fixed effects to increase precision.[12] We derive the

[12] As can be seen from Figure 1, the chance of experiencing violence for the first time is heavily dependent on the number of years in marriage.

predicted values for the conviction rate from the following first-stage regression

$$\text{ConvictRate}_{d,t-1} \;=\; \pi\, Z_{d,t-1} \;+\; \delta_d \;+\; \lambda_t \;+\; \nu_{dt}, \tag{5}$$

where $Z_{d,t-1}$ is the district–year instrument, constructed as described above. Note that we include district and year fixed effects (not district–year fixed effects); otherwise the district–year mean would absorb $\widehat{\text{ConvictRate}}_{dt}$, which varies only at this level. Standard errors are clustered at the level of districts.

4.3 Assessing the instrument

Before we proceed, we assess the instrument along several dimensions.

Validity. Validity of the instrument requires that judges are randomly assigned to cases, conditional on court, time and charge. In Section 2 we argue that this is plausible in our setting. In this section we test empirically whether our data is consistent with the quasi-random assumption.

We start by regressing a binary variable denoting conviction in the current case on court-year and charge section fixed effects and a set of pre-determined case characteristics: a female and Muslim dummy for the defendant, the defendant's lawyer, the petitioner and the petitioner's lawyer. Several of these characteristics have individual prediction power on conviction and the joint F statistics is highly significant, with a p-value<0.001, as can be seen from Column 1 of Table 1. This shows that the baseline covariates are strongly correlated with outcomes. In Column 2, we show that the same set of variables are *uncorrelated* with judge harshness, as we would expect if cases are allocated to judges in a random fashion. None of the variables are individually significant and the joint F statistics is far from being significant (p-value=0.610). We interpret these results as consistent with quasi-random assignment of court cases.

In the remainder of the table, we repeat the exercise at the district level. Remember that we cannot retain charge section fixed effects at this level of aggregation. Because of this, we also check whether the share of cases by charge sections predict the outcomes. The dependent variable in Column 3 is the district–year conviction rate in crimes against women, regressed on district and year fixed effects (as in our main specification), section shares and the set of case characteristics from above, now aggregated to the level of district-year. The section share controls are the only ones statistically significant at conventional levels, yielding a highly significant joint F-statistic of 3.7 ($p < 0.001$). In Column 4, we regress the judge–mix instrument on the same covariates. None of the individual coefficients are significant at the

5% level, and the joint F-test is also insignificant with a p-value of 0.281. This shows that the set of controls have very limited predictive content for the instrument, also after the aggregation to districts.

Relevance. Relevance requires that our instrument actually predicts convictions. We first assess this at the case level. To do this, we regress conviction in the current case on the harshness measure, controlling for court-year and charge section fixed effects. Estimates are shown in Column 1 of Table 2 and implies that a one–percentage–point increase in a judge's prior conviction rate is associated with a 0.33 percentage point higher chance of conviction in the current case. The effect is significant at a 1% level.

We next examine whether the judge mix measure predicts conviction rates at the district-year level. This is essentially our first-stage (effectively re-weighted according to the number of survey respondents in each district-year for the survey based outcome). Results are presented in Column 2. In this case, the estimates imply that a one–percentage–point increase in the average prior conviction rate is associated with a 0.21 percentage point higher conviction rate in the current year. The effect is again significant at the 1% level.

In Figure 3, we present a binned scatter plot, and in Figure 4 we visualize the variation in the instrument through a histogram plot. A large mass of the observations are centered around zero. Notably, around 40 percent of the district-year observations have an average prior conviction rate of exactly zero, meaning that none of the judges had any prior convictions. As shown in Appendix B, our main results hold also when excluding this zero subset (see Table A7).

Taken together, the results presented in this section show that conviction rates of crimes against women is strongly predicted by the type of judges a district ends up with in a particular year, i.e. that our instrument is relevant.

Exclusion. The exclusion restriction requires that the district–year judge mix affects violence only through its impact on convictions in crimes against women.

A natural threat is that stricter judges proxy broader, time-varying district changes — for instance prosecutorial drives, policing campaigns, or court-wide punitiveness — that could themselves affect violence. Two points help mitigate this concern. First, the instrument is constructed from pre-decision behavior and residualized for district–year shocks, weakening any mechanical link to contemporaneous district-level unobservables. Second, if the instrument merely proxied general prosecutorial vigor or generic courtroom "harshness", it should also raise convictions in other type of cases. Table A1 shows that this is not the case: the instrument has no prediction power of convictions in other types of crimes.

Monotonicity. If the effect of conviction on domestic violence is homogeneous, then quasi-random judge assignment and the exclusion restriction are sufficient conditions for 2SLS to recover the average causal effect. With heterogeneous effects, the standard LATE interpretation also requires *monotonicity*. In this case, monotonicity means that moving from a lenient to a harsher judge should weakly raise all cases' probability of conviction, i.e., there should be no cases that would be convicted by a lenient judge but acquitted by a harsher one. At the district level of aggregation, this implies that a harsher judge mix should weakly increase the conviction rate for all underlying case compositions.

An empirical implication of monotonicity is sign stability of the first stage across subsamples. We test this by re-estimating the first stage with the baseline instrument on a series of restricted samples. First, we partition the data by offence type: for each of the three most common IPC sections in our data, we run the regression separately on cases that do fall under that section and on those that do not. Second, we create subsamples based on whether the cases were decided in even versus odd calendar months. Third, we partion the sample based on case delays: within each district–year, we split cases by above-median and below-median delays. Under monotonicity, the instrument's coefficient should retain the same positive sign in all of these subsamples. This is the case, as can be seen from Panel A of Table A2.

In Panel B of Table A2 we implement the "reverse instrument" test of Bhuller et al. (2020). The intuition of this test is that a stricter judge in one context should also exhibit a higher propensity to convict in other contexts. Concretely, for each partition used in Panel A, we implement the test by first estimate the judge–harshness instrument in one subsample and then use this instrument for the opposite sub-sample. For example, in the even/odd month split, we use the instrument estimated from prior *even*-month cases to predict convictions in *odd*-month cases. As can be seen from the table, all the first stage estimates are positive, consistent with monotonicity.

5 Divergent effects of convictions on future violence

In this section we present our main results on the effect of convictions on violence against women. We document a clear divergence across data sources: higher conviction rates reduce survey-reported violence, while simultaneously increasing administrative records of crimes against women.

5.1 Empirical evidence

Table 3 reports our baseline estimates. Panel A is based on the survey-based measure of violence onset constructed from the household survey. Column (1) presents OLS estimates, and shows that a one–percentage-point increase in the conviction rate is associated with a 0.015 percentage-point decline in the probability of violence in the subsequent year. This estimate is statistically significant at the 10% level. Column (2) reports the reduced-form relationship between lagged judge harshness and the probability of violence. A one–percentage-point increase in judge harshness reduces the likelihood of violence by 0.049 percentage points, corresponding to about 3% of the sample mean and the effect is statistically significant at the 1% level. Column (3) presents 2SLS estimates instrumenting district-level conviction rates with judge harshness. The point estimate indicates that a one–percentage-point increase in the conviction rate in the previous year reduces the probability of violence by 0.224 percentage points, or about 15% of the sample mean. This effect is statistically significant at the 1% level.

Panel B uses the district-level number of reported crimes against women as outcome. The OLS estimates in Column (1) suggest that higher conviction rates are associated with less reported crimes: a one–percentage-point increase in the conviction rate in the previous year is associated with a 0.4 percentage-point reduction in reported crimes. In contrast, once we exploit variation from judge harshness, the pattern changes. The reduced-form estimates in Column (2) imply an increase of 0.6 percentage points, statistically significant at the 5% level; and the 2SLS estimates in Column (3) imply an increase of 2.8 percentage points, statistically significant at the 10% level.

Figure 5 displays dynamic effects for both outcomes, obtained by re-estimating the baseline 2SLS specifications with alternative leads and lags of the treatment variable. Each point corresponds to an estimate from a separate regression; in particular, the lag 1 estimate coincides with our baseline effect. The lead coefficients for the survey outcome are slightly positive but statistically insignificant, providing no evidence of differential pre-trends. The only statistically significant effect occurs one year after the decision year. Estimates at longer horizons (lag 2 and lag 3) are close to zero, suggesting that the effects on violence and reporting are short-lived.

Figure 6 shows that the effect on the survey violence measure is driven primarily by younger women. This is reassuring in light of Figure 1, which indicates that violence onset is substantially more likely early in marriage. We present a series of other robustness checks in Appendix B. We show, first, that the deterrence effect is strongest for women early in their marriage (consistent with the pattern in Figure 6). Second, we show that our results are robust to using more stringent measures of judge harshness, constructed by residualizing judges' decisions with either court–year–charge-section fixed effects or court–year–month–charge-section fixed effects. Third, we verify that the estimates are not driven by outliers. Fourth, we demonstrate that our findings are robust to restricting the sample to cases with the same filing and decision year, and to excluding district–years with no prior convictions for crimes against women.

In sum, our results point to a divergent response to more punitive courts: in the survey data, women are less likely to state that they have experienced violence, while in the administrative data, formally recorded crimes against women increase.

5.2 A simple model of deterrence and reporting

To interpret these divergent effects, we develop a simple framework in which an increase in conviction probabilities can simultaneously (i) deter violence and (ii) increase formally recorded crimes by raising victims' willingness to report.

Environment and timing. A couple consists of a potential perpetrator (husband) H and a victim (wife) W. A single period proceeds as follows:

1. The husband chooses the severity of violence $s \geq 0$.[13]
2. Observing s, the wife chooses whether to report to the police.
3. If she reports, the case is escalated with probability $\pi(s) \in [0,1]$, increasing in s. Conditional on escalation, the husband is convicted with probability $\mu \in [0,1]$ and receives sanction $F > 0$. We interpret μ as the model analogue of the conviction rate used in the empirical analysis, capturing the probability that an escalated case results in conviction.

Define the composite probability that a report leads to conviction as

$$q(s) \;\equiv\; \mu\, \pi(s). \tag{6}$$

Preferences. The husband's payoff is

$$U_H(s) \;=\; bs \;-\; \frac{c}{2}s^2 \;-\; r(s)\, q(s)\, F, \tag{7}$$

[13] Violence in intimate relationships may be partly strategic, but it is often also impulsive and triggered by momentary conflict. On way of interpreting the husband's choice of severity s is that it captures a self-chosen behavioral standard (a "rule") that shapes how he responds when situations escalate. Choosing a lower s can hence be understood as investing in self-discipline — adopting stricter internal norms that help keep spontaneous violence at a lower level. Under this interpretation, the model implies a strong correlation between the chosen norm s and realized violence.

where $b > 0$ captures the (reprehensible) benefit from using violence as a control device, $c > 0$ captures convex costs, and the final term is the expected legal cost. The wife's payoff from reporting is

$$U_W(s; r = 1) \;=\; -hs \;-\; \kappa \;+\; q(s)\,\phi, \tag{8}$$

while the payoff from not reporting is

$$U_W(s; r = 0) \;=\; -hs, \tag{9}$$

where $h > 0$ is harm from abuse, $\kappa \geq 0$ is the cost of reporting (stigma, time, fear), and $\phi \geq 0$ captures the benefit from successful prosecution (safety, validation, services). Let κ vary across wives with CDF $G(\cdot)$. Then the probability of reporting conditional on severity s is

$$r(s) \;\equiv\; \Pr(q(s)\phi \geq \kappa) \;=\; G(q(s)\phi) \;=\; G(\mu\,\pi(s)\phi)\,. \tag{10}$$

Anticipating (10), the husband chooses s to maximize

$$\max_{s \geq 0}\; bs - \frac{c}{2}s^2 - F\,\mu\,\pi(s)\,G(\mu\pi(s)\phi)\,, \tag{11}$$

which yields the following first-order condition

$$b - cs\,(\mu) - F\mu\,\pi_s(s\,(\mu))[G(\mu\pi(s\,(\mu))\,\phi) + \mu\phi\,\pi(s\,(\mu))\;G'(\mu\pi(s\,(\mu))\,\phi)] = 0. \tag{12}$$

Result 1 (Deterrence)**.** *The husband optimally reduces violence when the conviction probability μ is higher, as a higher μ increases the expected marginal cost of violence both by raising the chance that a given report leads to conviction and by making reporting more attractive.*

Formally, under $\pi_s(s) > 0$ and $G'(\cdot) \geq 0$ and standard regularity conditions ensuring a unique interior optimum, this implies the comparative static

$$\frac{ds\,(\mu)}{d\mu} < 0.$$

Administrative recorded crimes require that an incident is brought to the attention of the police and registered in the system. In our simple framework, two steps are essential: the wife must report, and the report must translate into an registered case. We capture the latter by the escalation function $\pi(s)$ and reporting by $r(s)$. Accordingly, we can model the probability of a recorded crime as[14]

$$C(\mu) \;\equiv\; \pi(s\,(\mu))\; r(s\,(\mu)). \tag{13}$$

Result 2 (Recorded crimes)**.** *Recorded crimes can rise even if actual violence falls, provided the reporting response dominates the deterrence-induced decline in incidents.*

Formally, an increase in μ lowers $s(\mu)$ (deterrence), which tends to reduce $\pi(s(\mu))$, but it also raises reporting incentives by increasing expected case success, which tends to increase $r(s(\mu))$. Hence the net effect of μ on $C(\mu)$ is a priori ambiguous.

Interpreting the empirical results. To connect the model more directly to the empirical specifications, it is convenient to work with logs. Taking logs of (13) yields

$$\log C(\mu) = \log \pi(s\,(\mu)) \;+\; \log r(s\,(\mu)). \tag{14}$$

Differentiating with respect to μ gives a decomposition of how conviction risk affects recorded crime:

$$\frac{d\log C(\mu)}{d\mu} = \underbrace{\frac{d\log \pi(s\,(\mu))}{d\mu}}_{\text{incidents}} + \underbrace{\frac{d\log r(s\,(\mu))}{d\mu}}_{\text{reporting}}. \tag{15}$$

This representation aligns with the administrative outcome, which we estimate in logs, and highlights that the semi-elasticity of recorded crimes with respect to conviction risk reflects two margins: an incidents component (through π) and a reporting component (through r). Empirically, the administrative coefficient $\widehat{\beta}_{\text{admin}}$ is therefore directly interpretable as a semi-elasticity of recorded crimes with respect to the conviction rate (measured in percentage points). By contrast, the survey outcome is in levels. To place it on a comparable scale, we convert the survey coefficient into an approximate semi-elasticity around the sample mean:

$$\widehat{\varepsilon}_V \;\equiv\; \frac{\widehat{\beta}_{\text{survey}}}{\bar{Y}_{\text{onset}}}. \tag{16}$$

We treat $\widehat{\varepsilon}_V$ as a proxy for the model's incidents response (the first term in (15)). This motivates the approximation

$$\widehat{\beta}_{\text{admin}} \;\approx\; \underbrace{\widehat{\varepsilon}_V}_{\text{incidents}} \;+\; \underbrace{\widehat{\gamma}_R}_{\text{reporting}}\,, \qquad \widehat{\gamma}_R \;\equiv\; \widehat{\beta}_{\text{admin}} - \widehat{\varepsilon}_V, \tag{17}$$

[14]This formulation abstracts from additional institutional steps between a report and an administrative record (e.g. FIR registration, classification, follow-through). These can be accommodated by allowing the report-to-record mapping to vary with severity and/or conviction risk, without changing the basic logic.

so that the gap between the administrative and survey responses can be interpreted as the implied reporting component needed to reconcile the two patterns.

In our data, $\widehat{\beta}_{\text{survey}} = -0.2240$ and $\bar{Y}^{\text{onset}} = 1.4862$, yielding $\widehat{\varepsilon}_V \approx -0.1507$. Interpreted as a semi-elasticity, this implies that a one–percentage-point increase in the conviction rate is associated with an approximately 15.1% decline in survey-based violence onset around the sample mean. With $\widehat{\beta}_{\text{admin}} = 0.0279$, the implied reporting semi-elasticity is $\widehat{\gamma}_R \approx 0.1786$, i.e. an approximately 17.9% increase in the reporting component per one–percentage-point increase in the conviction rate. Our empirical estimates are thus consistent with a setting in which conviction risk both discourages violence and reduces frictions to reporting — for instance, by increasing the perceived likelihood of case progression and sanctioning conditional on a report.

In Appendix C we provide a simple parametric version of the model that delivers closed-form expressions. The model example highlights how higher conviction risk can lead to lower actual violence and higher recorded crimes at the same time, through the joint operation of deterrence and reporting. Moreover, we also show that the magnitudes of our estimated effects can be rationalized within the model for plausible parameter values (i.e. when the reporting response dominates the deterrence-induced decline in incidents).

From enforcement to beliefs. The model treats μ as the relevant "conviction risk" faced by households. In practice, however, neither victims nor potential perpetrators observe conviction probabilities directly. Instead, they must form beliefs about the likelihood that a report will be escalated and ultimately result in conviction, and these beliefs may update only when enforcement becomes salient. This distinction matters because both margins in the model — deterrence (through the husband's choice of $s(\mu)$) and reporting (through $r(s) = G(\mu\pi(s)\phi)$) — operate through perceived case success.

Local media provides a natural channel for such belief updating. By publicizing court outcomes, discussing penalties and procedures, or circulating narratives about successful cases or the new law itself, local media can raise the salience of adjudication and improve legal literacy. Through the lens of the model, greater media penetration can therefore strengthen the behavioral response to a given change in underlying enforcement by (i) increasing the perceived conviction risk μ among potential perpetrators (amplifying deterrence) and (ii) lowering informational and social barriers to reporting (amplifying the reporting response). A direct empirical implication is that the deterrence effect of convictions on survey-based violence should be larger in high-media areas. At the same time, heterogeneity in recorded crimes is a priori less clear: if media simultaneously amplifies deterrence and reporting, the two effects may offset in the net change in $C(\mu) = \pi(s(\mu))\, r(s(\mu))$. Below we test these predictions by exploring whether the effects of court punitiveness are larger in area with greater exposure to local media.[15]

6 The role of local media

We examine whether media amplifies behavioral responses to enforcement using three complementary measures of local media exposure. First, we use self-reported media consumption from the NFHS. Second, we construct a plausibly more exogenous measure of radio access based on the geographic diffusion of community radio stations. Third, we proxy for newspaper reach using per-capita newspaper circulation at the district level. For each measure, we standardize exposure to mean zero and unit variance, interact it with our judge-harshness instrument, and estimate reduced-form specifications analogous to the baseline. We present 2SLS versions of the same specifications in Appendix E, where we instrument for both the conviction rate and its interaction with media exposure. The interaction term is generally weakly instrumented, so we place little weight on these estimates.

6.1 Measuring media exposure

Self-reported media use (NFHS). We use NFHS 2015–2016 questions on media usage, collected near the beginning of our study period. We classify respondents as regular consumers of newspapers, radio, or television if they report using the medium at least once per week. We then compute district-level shares of regular consumers for each medium and standardize these measures.

Radio access from community radio station locations. Radio remains one of the most accessible mass media in India, in part because it does not require literacy. We focus on community radio stations, which are designed to be local in both language and content production. The sector expanded following a 2006 policy reform that opened licenses to a limited set of organizations (educational institutions, NGOs, and government-financed agricultural centers) subject to centralized screening and licensing requirements. Programming rules emphasize development-oriented content in local languages or dialects and restrict certain types of broadcasts (notably political news). Content analyses suggest that women's empowerment is a central theme in many stations' programming (Pavarala and Malik, 2007; Rusche, 2025). These features make community radio a plausible channel for increasing legal

[15] A broader literature documents that media exposure can affect beliefs, attitudes, and socioeconomic outcomes in a variety of settings (see DellaVigna and La Ferrara, 2015, for a review).

literacy and lowering the perceived stigma or cost of reporting. We combine our main data with geocoded locations of community radio stations. For each NFHS survey cluster, we count the number of stations within a 50 km radius, following Rusche (2025), which captures approximately 96% of stations' total coverage area. We then standardize this exposure measure and interact it with judge harshness.

Newspaper access from district-level circulation. India also has an unusually local and vernacular newspaper market, with extensive district editions and regional-language dailies. To proxy for local print reach, we use the number of circulated newspapers per capita at the district level (measured in 2013–2014), standardize this measure and interact it with judge harshness as above.

6.2 Heterogeneity

Starting with the survey-based violence measure, the interaction coefficients for self-reported television and newspaper use are close to zero (Columns 1 and 3, Table 4), whereas the interaction with self-reported radio use is negative and statistically significant at the 5% level (Column 2, Table 4). The magnitude implies that a one-standard-deviation increase in radio exposure is associated with a 43% larger (i.e., more negative) effect of court punitiveness on violence.

Using the plausibly more exogenous proxy based on community radio station locations, we again find a negative and statistically significant interaction between judge harshness and radio access (Column 1, Table 5). A one-standard-deviation increase in community radio exposure makes the reduced-form effect of judge harshness on violence about 3 percentage points more negative, which is more than half of the average reduced-form effect. We also estimate a negative main effect of radio access on violence onset, consistent with community radio independently reducing domestic violence through information and norm-shifting channels, in line with Rusche (2025). We are able to include the main variable of media exposure in this specification (and not just the interaction) as we have time-varying information on the number of local radio stations at the level of survey clusters.

Finally, we obtain a similar pattern using district-level newspaper circulation (Column 1, Table 6): the interaction is negative, and the magnitude suggests that the deterrence effect becomes about 1.4 percentage points more negative with a one-standard-deviation increase in newspaper circulation.

For administrative crime counts, the heterogeneity results are notably weaker: the interaction terms between judge harshness and the different media exposure measures are all close to zero and far from being statistically significant at conventional levels.

Taken together, these results suggest that local media — especially radio — amplifies the deterrence response to changes in court punitiveness as measured by self-reported violence. At the same time, we do not detect comparable heterogeneity in recorded crimes. This contrast underscores that administrative records conflate two margins. If media simultaneously (i) reduces underlying violence through deterrence and (ii) lowers the cost or stigma of reporting by victims, then the net effect on recorded crime may be small. Through the lens of the model, one interpretation is that higher media penetration amplifies both deterrence and reporting, so that the additional effect of enforcement in high-media areas largely cancels out in the net change in recorded crime.

More broadly, our results suggest that adjudication reforms can have larger behavioral effects when complemented by channels that translate legal change into widely understood, locally salient information. This complementarity is likely to be particularly important in settings with low legal literacy and substantial social stigma surrounding reporting.

7 Conclusion

This paper provides new causal evidence that criminal convictions deter domestic violence in India. Using judge assignment as an instrument for district-year conviction rates, we show that higher conviction probabilities reduce the onset of intimate partner violence, particularly in areas with stronger local media. At the same time, reported crimes rise, consistent with increased formalization rather than higher victimization. The combination of declining self-reported violence and rising reports underscores the importance of moving beyond official crime statistics when evaluating enforcement.

In all, our evidence is most consistent with an information-driven deterrence mechanism that operates quickly through perceived sanction risk. The short–lived nature of the effect suggests a salience channel rather than deep norm change among men: convictions become salient, beliefs about the probability of punishment rise, and offending falls in the short run; as media attention fades and judges rotate, the effect dissipates. The result of increased number of reported crimes against women over the same horizon is consistent with updating beliefs among women — related to both the acceptability of beating and the chance of getting through the system with a crime report.[16]

The policy implications are twofold. First, strengthening the credibility and visibility of adjudication can yield immediate preventive benefits. Investments in court efficiency and media transparency may amplify these effects. Second, sustaining change requires repetition:

[16]In line with this, we find suggestive evidence of a decline in women's acceptance of wife beating, whereas men's attitudes are unchanged, see Appendix D.

lasting norm shifts likely depend on consistent enforcement and ongoing public dissemination of outcomes. These results suggest that justice, when both credible and visible, can serve as deterrence — but turning short-term deterrence into durable norm transformation remains a central challenge.

References

Aizer, A. and J. J. Doyle (2013): "Juvenile Incarceration, Human Capital and Future Crime: Evidence from Randomly-Assigned Judges," .

Amaral, S. (2017): "Do improved property rights decrease violence against women in India?" Working Paper 2017-13, ISER Working Paper Series.

Amaral, S., S. Bhalotra, and N. Prakash (2021): "Gender, Crime and Punishment: Evidence from Women Police Stations in India," .

Anderson, S. (2021): "Intimate partner violence and female property rights," *Nature Human Behaviour*, 5, 1021–1026, publisher: Nature Publishing Group.

Arnold, D., W. Dobbie, and C. S. Yang (2018): "Racial Bias in Bail Decisions*," *The Quarterly Journal of Economics*, 133, 1885–1932.

Ash, E., S. Asher, A. Bhowmick, S. Bhupatiraju, D. Chen, T. Devi, C. Goessmann, P. Novosad, and B. Siddiqi (2025): "In-group bias in the Indian judiciary: Evidence from 5 million criminal cases," *Review of Economics and Statistics*, 1–45.

Behrer, A. P., D. L. Chen, S. Joshi, O. Kyrychenko, V. Nagarathinam, P. Neis, and S. Singh (2025): "Decoding Green Justice: An AI-Assisted Exploration of Indian Environmental Rulings over Three Decades," .

Beleche, T. (2019): "Domestic violence laws and suicide in Mexico," *Review of Economics of the Household*, 17, 229–248.

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

Bhupatiraju, S., D. Chen, S. Joshi, P. Neis, and S. Singh (2024): "Environmental Litigation as Scrutiny: A Four Decade Analysis of Justice, Firms, and Pollution in India," .

Bobonis, G. J., M. González-Brenes, and R. Castro (2013): "Public Transfers and Domestic Violence: The Roles of Private Information and Spousal Control," *American Economic Journal: Economic Policy*, 5, 179–205.

Cage, J., G. Cassan, and F. R. Jensenius (2023): "Political Determinants of the News Market: Novel Data and Quasi-Experimental Evidence from India," .

Corradini, V. and G. Buccione (2023): "Unilateral divorce rights, domestic violence and women's agency: Evidence from the Egyptian *Khul* reform," *Journal of Development Economics*, 160, 102947.

Das, A. and C. M. Lakshmana (2020): "The Implementation of Domestic Violence Act in India: A State-Level Analysis," Working paper 499, Institute for Social and Economic Change, Bangalore.

DellaVigna, S. and E. La Ferrara (2015): "Economic and social impacts of the media," in *Handbook of media economics*, Elsevier, vol. 1, 723–768.

Do, Q.-T., S. Joshi, and S. Stolper (2018): "Can environmental policy reduce infant mortality? Evidence from the Ganga Pollution Cases," *Journal of Development Economics*, 133, 306–325.

Ferraz, C. and L. Schiavon (2022): "Crime, Punishment, and Prevention: The Effect of a Legal Reform on Violence Against Women," .

García-Ramos, A. (2021): "Divorce laws and intimate partner violence: Evidence from Mexico," *Journal of Development Economics*, 150, publisher: Elsevier.

Gulesci, S., M. Leone, and S. Zafar (2024): "Domestic Violence Laws and Social Norms: Evidence from Pakistan," *Trinity Economics Papers*, number: tep0324 Publisher: Trinity College Dublin, Department of Economics.

Harari, M. (2019): "Women's Inheritance Rights and Bargaining Power: Evidence from Kenya," *Economic Development and Cultural Change*, 68, 189–238, publisher: The University of Chicago Press.

Mehmood, S., A. Seror, and D. L. Chen (2023): "Ramadan fasting increases leniency in judges from Pakistan and India," *Nature Human Behaviour*, 7, 874–880, publisher: Nature Publishing Group.

Minster of Broadcasting and Information (2020): "State wise details of Operational Community Radio Stations in India as on 31-03-2020," https://web.archive.org/web/20200818181824/https://mib.gov.in/sites/default/files/List%20of%20289%20Operational%20CRS%20to%20upload%20on%20website%20PDF.pdf.

Pavarala, V. and K. K. Malik (2007): *Other Voices: The Struggle for Community Radio in India*, New Delhi: SAGE Publications India Pvt. Ltd.

Perova, E. and S. A. Reynolds (2017): "Women's police stations and intimate partner violence: Evidence from Brazil," *Social Science & Medicine*, 174, 188–196.

Ramos-Maqueda, M. and D. L. Chen (2025): "The data revolution in justice," *World Development*, 186, 106834.

Rusche, F. (2025): "Broadcasting change: India's community radio policy and women's empowerment," Tech. rep., Discussion Papers of the Max Planck Institute for Research on Collective Goods.

Sanin, D. (2025): "Women's Employment, Husbands' Economic Self-Interest and Domestic Violence," .

Shreemoyee, S., P. Roychowdhury, and G. Dhamija (2025): "Women's attitudes towards physical intimate partner violence in India: Trends, patterns, and determinants," *PloS one*, 20, e0318350.

Sviatschi, M. M. and I. Trako (2024): "Gender violence, enforcement, and human capital: Evidence from women's justice centers in Peru," *Journal of Development Economics*, 168, 103262.

UN Women (2014): "The Costs of Violence," Tech. rep.

UNDP (2022): "Human Development Report 2021-22," *UNDP (United Nations Development Programme).*

Vyas, P. (2006): "Reconceptualizing Domestic Violence in India: Economic Abuse and the Need for Broad Statutory Interpretation to Promote Women's Fundamental Rights," *Michigan Journal of Gender & Law*, 13, 177–206.

Tables

Table 1: Testing the validity of the instrument

	Case-level		District-level	
Dep.var.	Convicted	Judge harshness	Conviction rate	District harshness
	(1)	(2)	(3)	(4)
Religion defendant (=1 Muslim)	-0.090	-0.032	-0.604	0.049
	(0.095)	(0.032)	(1.943)	(0.739)
Gender defendant (=1 female)	-0.017	0.018	-1.568	-0.420
	(0.172)	(0.039)	(1.405)	(0.559)
Religion defendant's lawyer (=1 Muslim)	-0.794***	-0.150	0.380	-0.229
	(0.276)	(0.106)	(0.618)	(0.325)
Gender defendant's lawyer (=1 female)	1.388**	-0.139	1.824**	1.027*
	(0.567)	(0.099)	(0.880)	(0.526)
Religion petitioner (=1 Muslim)	0.261**	0.049	-0.230	1.042
	(0.101)	(0.037)	(0.922)	(0.684)
Gender petitioner (=1 female)	0.084	0.078	0.552	-0.533
	(0.098)	(0.055)	(0.486)	(0.370)
Religion petitioner's lawyer (=1 Muslim)	0.028	-0.059	-0.801	-0.354
	(0.130)	(0.069)	(0.881)	(0.439)
Gender petitioner's lawyer (=1 female)	-0.514	0.092	0.759	0.520
	(0.330)	(0.140)	(0.825)	(0.469)
Charge-section share (IPC 312)			-2.835	-3.607*
			(3.594)	(2.113)
Charge-section share (IPC 354)			5.563***	0.707
			(1.534)	(0.838)
Charge-section share (IPC 376)			6.728***	1.115
			(1.587)	(1.000)
Charge-section share (IPC 498)			-0.829	-0.387
			(1.075)	(0.591)
Observations	172965	172965	1882	1882
Dep.var.mean	1.392	0.037	3.732	0.204
Clusters	472	472	436	436
Joint F-stat	4.925	0.866	2.78	1.168
[p-value]	[0.000]	[0.610]	[0.000]	[0.281]

Note: Columns (1)–(2) report case-level regressions with court-by-year and legal-section fixed effects. Column (1) shows that predetermined case characteristics predict convictions; Column (2) shows that the same characteristics do not predict judge harshness, consistent with quasi-random case assignment. Columns (3)–(4) repeat analogous balance checks at the district-year level (with district and year fixed effects); section shares are included where indicated to account for charge composition. Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 2: Relevance of instrument: Judge harshness and convictions

	Case-level	**District-level**
	(1)	(2)
Judge Harshness	0.330***	0.211***
	(0.038)	(0.059)
Observations	172965	1882
Dep.var.mean	1.392	3.732
Clusters	472	436

Note: The regressions in the table test the relevance of the instrument. The unit of observation is court cases in Column 1 and district-year in Column 2. Column 1 includes court-year and charge section fixed effects, while Column 2 includes district and year fixed effects. "Dep.var.mean" displays the average conviction rate. Standard errors clustered on districts are shown in the parentheses. ***p<0.01, **p<0.05, *p<0.10.

Table 3: Main results: Convictions, judge harshness and violence

	OLS	Reduced-form	2SLS
	(1)	(2)	(3)
Panel A: Violence onset (survey)			
Share convictions t-1	-0.015*		-0.224***
	(0.009)		(0.084)
District harshness t-1		-0.049***	
		(0.016)	
Observations	102614	102614	102614
Dep.var.mean	1.486	1.486	1.486
F-stat excl. instrument			17.9
Clusters	384	384	384
Panel B: Log reported crimes (admin data)			
Share convictions t-1	-0.004*		0.028*
	(0.002)		(0.016)
District harshness t-1		0.006**	
		(0.002)	
Observations	1875	1875	1875
Dep.var.mean	3.324	3.324	3.324
F-stat excl. instrument			12.7
Clusters	435	435	435

Note: The table presents our main estimates. Administrative outcomes are at the district-year level; survey outcomes are at the individual level and include marriage-duration fixed effects. Reduced-form columns regress outcomes on district harshness; 2SLS columns instrument the conviction rate with district harshness. All regressions include district and year fixed effects. Standard errors are clustered by district. "Dep.var.mean" displays the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: Self-reported media use, reduced-form estimates

	Violence onset			Reported crimes		
	News-paper	Radio	TV	News-paper	Radio	TV
	(1)	(2)	(3)	(4)	(5)	(6)
District harshness t-1	-0.050***	-0.049***	-0.050***	0.006**	0.006**	0.006***
	(0.015)	(0.015)	(0.018)	(0.002)	(0.002)	(0.002)
District harshness t-1 ×Media use	-0.005	-0.021***	0.003	-0.003	-0.000	-0.003
	(0.017)	(0.008)	(0.022)	(0.002)	(0.001)	(0.002)
Observations	102614	102614	102614	1875	1875	1875
Dep.var.mean	1.486	1.486	1.486	3.324	3.324	3.324
Clusters	384	384	384	435	435	435
Media, level	district	district	district	district	district	district

Note: The table tests for heterogeneity by local media exposure at the level of districts. Each column reports a reduced-form regression (outcome on district harshness) augmented with an interaction between district harshness and the indicated media measure, as described in the text. The specification otherwise follows the baseline specification (district and year fixed effects; survey regressions also include marriage-duration fixed effects). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Access to community radio states and violence, reduced-form estimates

	Violence onset	Reported crimes
	(1)	(2)
District harshness t-1	-0.050***	0.005**
	(0.016)	(0.002)
Radio	-0.163	0.018
	(0.101)	(0.015)
District harshness t-1 ×Radio	-0.033***	-0.001
	(0.012)	(0.002)
Observations	102614	1623
Dep.var.mean	1.486	3.343
Clusters	384	370
Radio, level	clusters	district

Note: The table tests for heterogeneity by community radio availability at the level of survey clusters/districts. Reduced-form regressions interact district harshness with an indicator for access to community radio (number of stations within 50 km). The regressions otherwise follow the baseline specifications (district and year fixed effects; survey regressions also include marriage-duration fixed effects). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6: Access to local newspapers and violence, reduced-form estimates

	Violence onset (1)	**Reported crimes** (2)
District harshness t-1	-0.064***	0.005**
	(0.015)	(0.002)
District harshness t-1	-0.014**	0.001
× Newspaper circulation p.c. (std)	(0.007)	(0.001)
Observations	94127	1647
Dep.var.mean	1.409	3.338
Clusters	380	377
Newspaper, level	district	district

Note: The table test for heterogeneity by local newspaper circulation at the level of districts. Reduced-form regressions interact district harshness with the district-level measure of newspaper circulation per capita (in 2013-2014). The regressions otherwise follow the baseline specifications (district and year fixed effects; survey regressions also include marriage-duration fixed effects). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figures

Figure 1: Risk of first violence, by years in marriage.

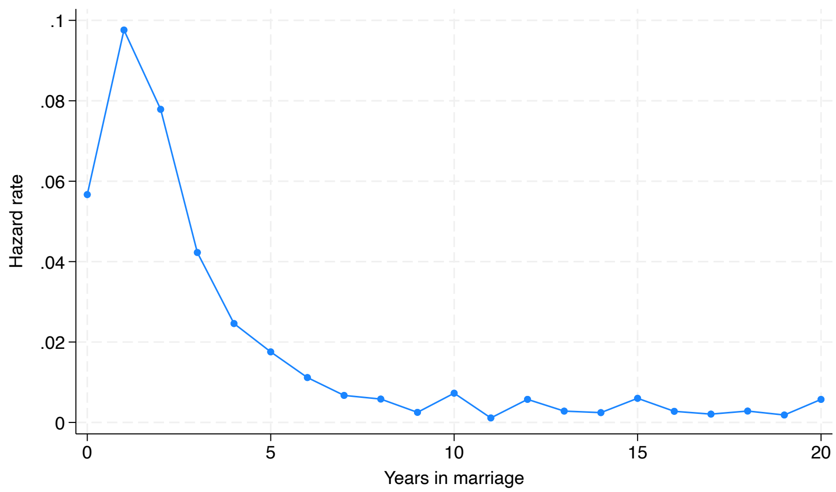

Note: The figure shows the hazard rate for women experiencing violence for the first time in a given marriage year, divided by the number of women at risk in that marriage year (i.e. those that have not yet experienced violence).

Figure 2: Community Radio Stations as of 2020

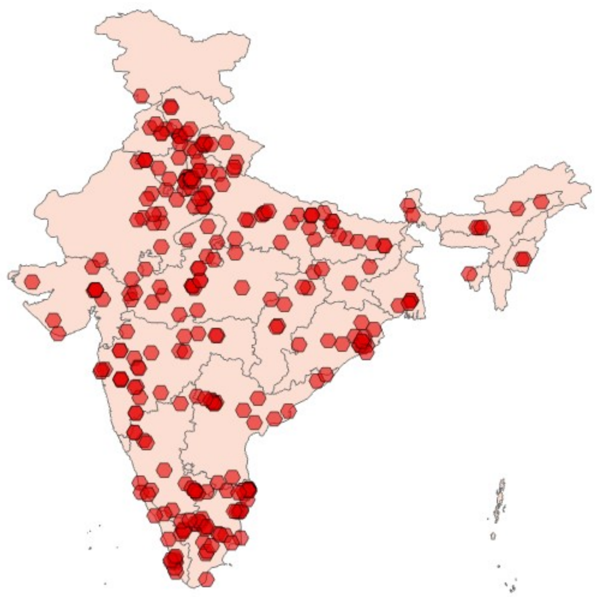

Note: The figure shows a map of community radio stations as of 2020. Station locations are geocoded from address information and overlaid on district boundaries (see data section for details and sources).

Figure 3: Binscatter: Conviction rates and predicted judge harshness

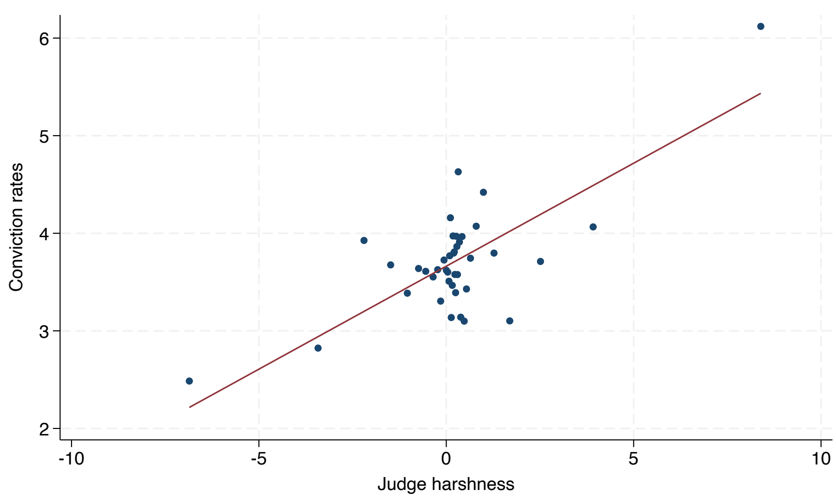

Note: The figure shows a binscatter of district-year conviction rates against predicted judge harshness, using 40 equally sized bins.

Figure 4: Variation in the instrument

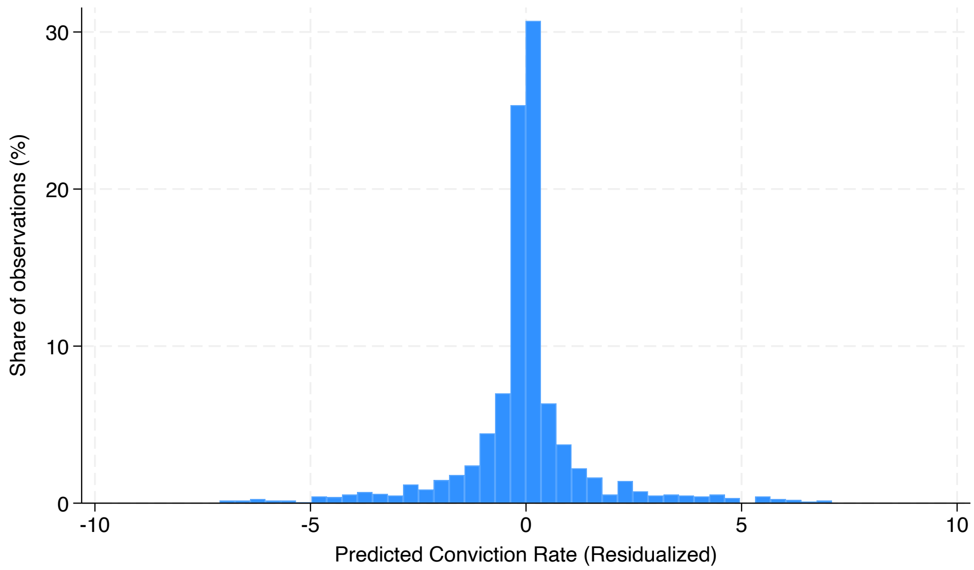

Note: The figure shows the distribution of the residualized predicted conviction rate (constructed from the judge-harshness instrument), after removing fixed effects and controls as in the baseline first-stage regression.

Figure 5: Dynamic effects on violence onset and reported crimes

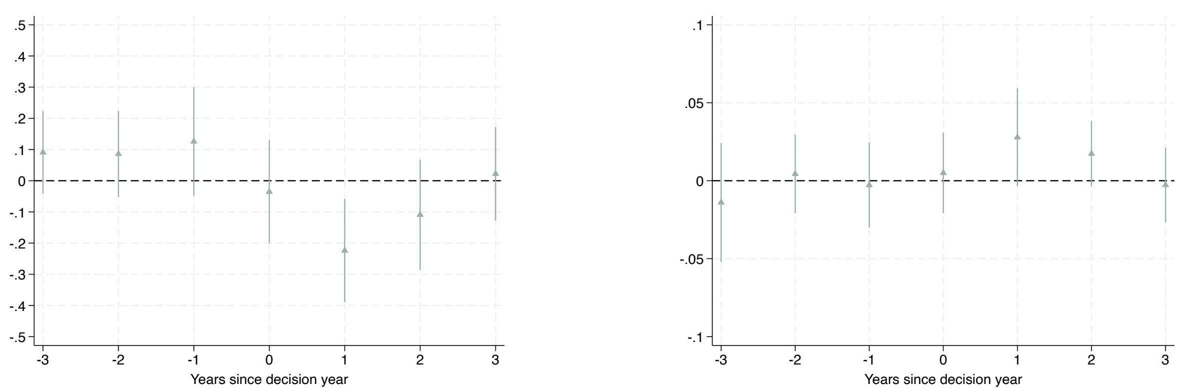

(a) Violence onset **(b)** Reported crimes

Note: The figure shows estimates for different years around the decision year ($t = 0$), shown for (a) survey-based violence onset and (b) log reported crimes. Each point comes from a separate regression for the corresponding event time; bars denote 95% confidence intervals. Standard errors are clustered by district.

Figure 6: Heterogeneous effects on violence onset, by age groups

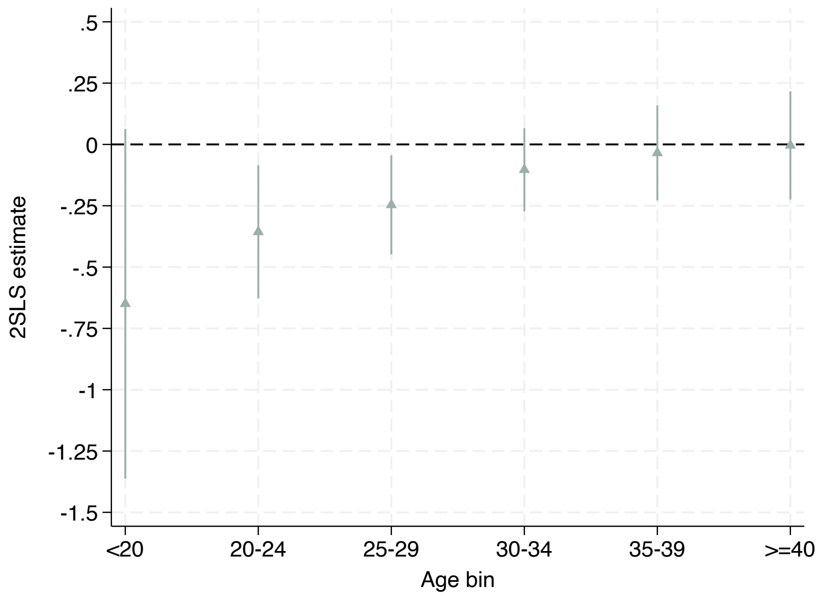

Note: The figure shows heterogeneous 2SLS estimates by age group. Each point reports the estimated effect for the indicated age bin; bars denote 95% confidence intervals. Standard errors are clustered by district.

Online Appendix

A Validation checks

Counterfactual random assignment of cases. Judges might move between courts, raising the concern that variation in our instrument could be driven by changes in which judges are available rather than by quasi-random case allocation. As a placebo, we therefore simulate counterfactual random assignment of cases to the set of judges available in each court-year. We define the available set as all judges observed in a given court in that year in any type of case who have previously handled at least one crimes-against-women case. For each simulated assignment we re-estimate the judge-harshness measure and then re-run the first-stage regression. Repeating this procedure for 1,000 random draws yields a placebo distribution of first-stage coefficients that is centered near zero, though slightly right-skewed, see Figure A1. In contrast, the first-stage coefficient obtained under the actual case assignment lies far in the tail of this placebo distribution. In fact, none of the 1,000 placebo draws produced an estimate as large (in absolute value) as the one observed in the data. This suggests quite strongly that the identifying variation in our instrument arising from realized case allocation rather than from changes in judge availability over time.

Judge harshness and other types of crimes. The exclusion restriction requires that the district–year judge mix affects violence only through its impact on convictions in crimes against women. A natural threat is that stricter judges proxy broader, time-varying district changes — for instance prosecutorial drives, policing campaigns, or court-wide punitiveness — that could themselves affect violence. In Figure A1 we therefore test whether our judge-harshness instrument predicts conviction rates in crimes other than crimes against women. If the instrument merely captures general prosecutorial vigor or generic courtroom "harshness," we would expect it to predict convictions across a broader set of offences. However, as shown in the table, the judge-harshness measure — constructed solely from earlier behavior in crimes-against-women cases — has no predictive power for convictions in other crime categories, supporting the interpretation that the instrument isolates variation specific to adjudication in crimes against women.

Monotonicity. Table A2 presents the monotonicity tests described in detail in Subsection 4.3 of the main paper. With heterogeneous treatment effects, the standard LATE interpretation of 2SLS additionally requires monotonicity: moving from a more lenient to a harsher judge should weakly increase (and never decrease) the probability of conviction for any given case. A direct empirical implication is sign stability of the first stage across subsamples. In Table A2 Panel A, we therefore re-estimate the first stage on a set of restricted samples (by offence type, even vs. odd months, and above- vs. below-median case delays) and find that the instrument retains the same positive sign throughout. Panel B implements the "reverse instrument" test of Bhuller et al. (2020); across all partitions, harshness estimated in one subsample predicts convictions in the complementary subsample with a positive coefficient. Together, these results are consistent with monotonicity.

Figure A1: Counterfactual random assignment of cases, distribution of first-stage estimates

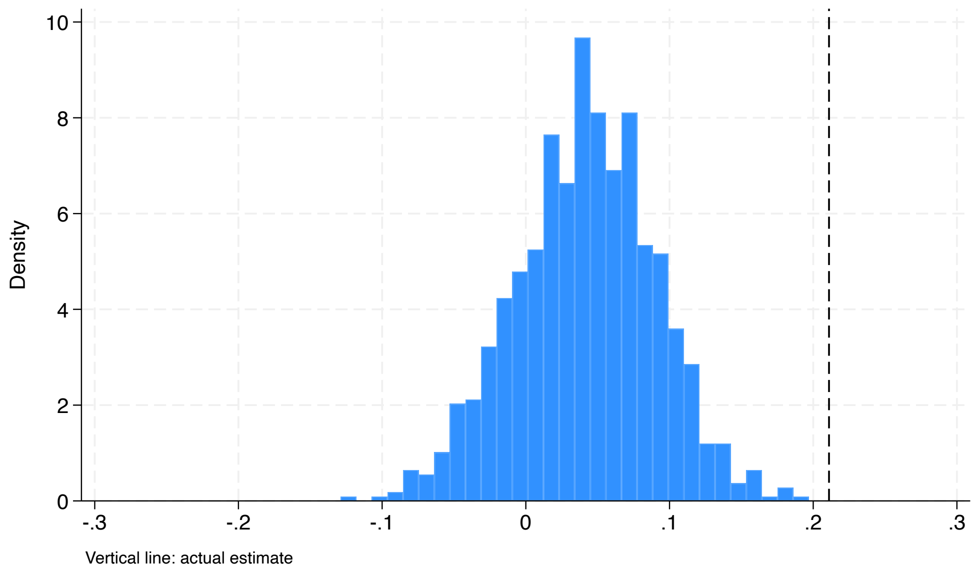

Note: The figure plots the distribution of first-stage coefficients from 1,000 placebo simulations in which, within each court-by-year cell, cases are randomly reassigned to the judges observed in that cell (thereby holding fixed the set of "available" judges and the number/composition of cases). For each draw we re-estimate the judge-harshness measure using the simulated case allocation and re-run the baseline first-stage regression. The vertical line shows the first-stage coefficient from the actual data.

Table A1: Judge harshness and convictions in other types of cases

Dep. var.	All other cases (1)	Property crime (2)	Personal crime (3)	Other crimes (4)
Judge Harshness	-0.003	0.109*	0.093	-0.043
	(0.057)	(0.064)	(0.079)	(0.082)
Observations	1882	1882	1882	1880
Dep.var.mean	5.781	5.334	8.282	9.663
Clusters	436	436	436	436

Note: The unit of observation in the table is district-year. The dependent variables in the table are average conviction rates for all crimes not defined as crime against women (Column 1), property crimes (Column 2), personal crimes (with a male defendant only) (Column 3), and other types of crimes (Column 4). "Dep.var.mean" displays the average conviction rate. Standard errors clustered on districts are shown in the parentheses. ***p<0.01, **p<0.05, *p<0.10.

Table A2: Testing Monotonicity

	IPC 366		IPC 498		IPC 354		Even month		Above median delay	
	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Baseline instrument										
Case level										
Judge Harshness	0.380***	0.188***	0.340***	0.071	0.252***	0.559***	0.298***	0.357***	0.450***	0.227***
	(0.048)	(0.041)	(0.041)	(0.054)	(0.037)	(0.066)	(0.041)	(0.053)	(0.062)	(0.063)
Observations	117174	55268	115608	56677	138227	34006	77555	88626	50906	61432
Dep.var.mean	.014	.014	.019	.004	.011	.024	.014	.014	.035	.009
Clusters	466	434	466	407	466	449	463	466	437	455
District level										
District Harshness	0.206***	0.313***	0.201**	0.054	0.242***	0.140	0.207***	0.211**	0.183	0.200***
	(0.070)	(0.108)	(0.084)	(0.094)	(0.084)	(0.097)	(0.071)	(0.088)	(0.125)	(0.060)
Observations	1881	1856	1882	1847	1881	1860	1881	1879	1880	1881
Dep.var.mean	.038	.037	.046	.016	.033	.049	.037	.039	.057	.024
Clusters	436	434	436	433	436	434	436	436	436	436
Panel B: Reverse sample instrument										
Case level										
Judge Harshness	0.127	0.116***	0.273**	0.027	0.083***	0.436***	0.202***	0.230***	0.280**	0.086*
	(0.083)	(0.043)	(0.111)	(0.019)	(0.028)	(0.102)	(0.048)	(0.042)	(0.130)	(0.045)
Observations	53688	46161	39868	42825	71748	29516	66972	73864	26112	49413
Dep.var.mean	.015	.013	.009	.003	.008	.024	.013	.014	.032	.009
Clusters	387	413	346	367	383	434	439	434	341	400
District level										
District Harshness	0.070	0.181	0.054	0.012	0.131***	0.111	0.082*	0.016	0.195**	0.095
	(0.062)	(0.135)	(0.153)	(0.080)	(0.044)	(0.085)	(0.048)	(0.077)	(0.081)	(0.066)
Observations	1356	1744	1282	1691	1224	1790	1670	1656	1293	1676
Dep.var.mean	.038	.037	.041	.015	.03	.049	.034	.038	.058	.024
Clusters	356	420	330	412	329	424	409	402	352	395

Note: With heterogeneous treatment effects, the standard LATE interpretation of our 2SLS estimates additionally requires monotonicity: assignment to a harsher judge should weakly increase each case's probability of conviction, so that a harsher judge mix weakly raises district-level conviction rates for any underlying case composition. A testable implication is sign stability of the first stage across subsamples. We assess this by re-estimating the first stage within partitions by offence type (top IPC sections vs. other cases), decision timing (even vs. odd months), and case delays (above- vs. below-median within district–year). In all partitions, the instrument retains a positive sign (Panel A). We also implement the "reverse instrument" test of Bhuller et al. (2020), constructing judge harshness in one subsample and applying it to the complementary subsample. For example, in the even/odd month split, we use the instrument estimated from prior *even*-month cases to predict convictions in *odd*-month cases. As can be seen from Panel B, the first-stage coefficients remain positive throughout, consistent with monotonicity.
"Mean" displays the average values of the dependent variables. Standard errors clustered on districts are shown in the parentheses.
***p<0.01, **p<0.05, *p<0.10.

B Robustness checks

In this section we present a series of robustness tests of our main results.

Restricted to women early in their marriage. Figure 1 indicates that violence typically begins within the first few years of marriage and rarely begins more than ten years after marriage. The figure also suggests some heaping in reported time-to-onset at round numbers (10, 15, and 20 years). As a robustness check, we therefore re-estimate our main specifications restricting the sample to women who are at most ten or five years into marriage at the time of the survey (see Table A3).

The IV point estimates are larger in absolute magnitude in these restricted samples, implying reductions of 0.52 percentage points (max 10 years in marriage) and 0.94 percentage points (max 5 years in marriage). However, baseline violence is also higher in these subsamples. Scaled by the corresponding sample means, the implied effects are around 13%, which is close to the baseline semi-elasticity of about 15%.

Alternative measures of judge harshness. As a second robustness check, we use more stringent specification when estimating case level judge harshness. Specifically, we replace the district-year effects in Equation 1 with either district-court-year fixed effects or district-court-year-legal section fixed effects. As shown in Table A4, the estimates barely change when using these specifications, but they lose some statistic significance.

Outliers. We next investigate whether our results are driven by sample outliers. To do this, we winsorize both the judge harshness measure and the district-level conviction rates. In Table A5 we show 2SLS estimate when we winsorize at 1% and 99% (Columns 1-2), and at 5% and 95% (Columns 3-4). Our results remain robust in both of specifications.

Identical decision and filing year. We construct our district-level judge-harshness measure by aggregating judge–year harshness to the district and *decision*-year level (see Subsection 4.1). A potential concern is that cases are quasi-randomly assigned to judges at the time of *filing*, whereas our aggregation uses decision year. As a robustness check, we therefore restrict the sample to cases with identical filing and decision years. Table A6 reports the resulting estimates. The point estimates are very similar to our baseline findings, suggesting that this timing mismatch does not drive our conclusions; statistical significance is however weaker due to the smaller sample size.

Excluding district-years with zero convictions. A large mass of the the district-year observations have an average prior conviction rate of exactly zero, meaning that none of the judges had any prior convictions (see Figure 4). To examine whether these observations are driving our results, we run the main specifications excluding this zero subset. As can be seen from Table A7, the estimates change very little as compare to the baseline. However, we lose some statistical significance due to the smaller sample sizes.

Table A3: Robustness: Restricted to women early in their marriage

	No more than 10y in marriage		**No more than 5y in marriage**	
	Reduced-form	2SLS	Reduced-form	2SLS
	(1)	(2)	(3)	(4)
Share convictions t-1		-0.521**		-0.938*
		(0.243)	(0.492)	
District harshness t-1	-0.118**		-0.197*	
	(0.048)		(0.107)	
Observations	30664	30664	9694	9694
Dep.var.mean	4.148	4.148	7.066	7.066
F-stat excl. instrument		18.1		10.6
Clusters	384	384	382	382

Note: The regressions in the table are restricted to women early in marriage. The sample used in Columns 1-2 is restricted to women who are at most 10 years into marriage, while the sample in Columns 3-4 is restricted to women at most 5 years into marriage. The regressions otherwise follow the baseline specifications with district and year fixed effects. Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A4: Robustness: Alternative judge harshness measure

	Court-year-section FE		**Court-year-section-month FE**	
	Reduced-form	2SLS	Reduced-form	2SLS
	(1)	(2)	(3)	(4)
Panel A: Violence onset (survey)				
Share convictions t-1		-0.196**		-0.276**
		(0.091)		(0.127)
District harshness t-1	-0.048**		-0.045***	
	(0.020)		(0.015)	
Observations	102570	102570	98865	98865
DepVarMean	1.487	1.487	1.48	1.48
F-stat excl. instrument		12.7	6.7	
Clusters	384	384	378	378
Panel B: Log reported crimes (admin data)				
Share convictions t-1		0.024		0.024
		(0.018)		(0.024)
District harshness t-1	0.005*		0.003	
	(0.003)		(0.003)	
Observations	1874	1874	1818	1818
Dep.var.mean	3.324	3.324	3.335	3.335
F-stat excl. instrument		8.1		3.9
Clusters	435	435	429	429

Note: Relative to the baseline harshness estimation (which conditions on district-year effects when estimating judge-year harshness), this table recomputes case-level harshness using stricter fixed effects when estimating harshness: district-court-year-legal-section fixed effects and, in the most restrictive specification, district-court-year-legal-section-month fixed effects (as indicated in the column headers). Reduced-form and 2SLS estimates are then re-run using these alternative instruments. The remaining regression structure follows the baseline specification (district and year fixed effects; survey regressions include marriage-duration fixed effects as well). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A5: Robustness: Winsorizing

	[1%, 99%]		**[5%, 95%]**	
	Reduced-form (1)	2SLS (2)	Reduced-form (3)	2SLS (4)
Panel A: Violence onset (survey)				
Share convictions t-1		-0.244**		-0.303**
		(0.101)		(0.145)
District harshness t-1	-0.059***		-0.078**	
	(0.021)		(0.032)	
Observations	102614	102614	102614	102614
DepVarMean	1.486	1.486	1.486	1.486
F-stat excl. instrument		12.8		8.4
Clusters	384	384	384	384
Panel B: Log reported crimes (admin data)				
Share convictions t-1		0.024		0.047*
		(0.018)		(0.025)
District harshness t-1	0.007**		0.011**	
	(0.003)		(0.004)	
Observations	1875	1874	1875	1875
Dep.var.mean	3.324	3.324	3.324	3.324
F-stat excl. instrument		8.1		10.6
Clusters	435	435	435	435

Note: Relative to the baseline, this table winsorizes the judge-harshness measure and the district-level conviction rate at the indicated cutoffs (1%/99% and 5%/95%) and then re-estimates the reduced-form and 2SLS specifications. All other elements of the specification follow the baseline (district and year fixed effects; survey regressions include marriage-duration fixed effects as well). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A6: Robustness: filing year=decision year

	Violence onset (survey)		Reported crimes (admin)	
	Reduced-form	2SLS	Reduced-form	2SLS
	(1)	(2)	(3)	(4)
Share convictions t-1		-0.258*		0.023
		(0.155)		(0.025)
District harshness t-1	-0.036***		0.002	
	(0.011)		(0.001)	
Observations	75402	75402	1443	1443
Dep.var.mean	1.442	1.442	3.358	3.358
F-stat excl. instrument		4		2
Clusters	311	311	368	368

Note: Relative to the baseline, the sample is restricted to cases with the same filing year and decision year to address potential mismatch between the timing of quasi-random assignment (filing) and the timing used for aggregation (decision year). Reduced-form and 2SLS specifications are otherwise identical to the baseline (district and year fixed effects; baseline controls; survey regressions include marriage-duration fixed effects as well). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A7: Robustness: Excluding zero

	Violence onset (survey)		Reported crimes (admin)	
	Reduced-form	2SLS	Reduced-form	2SLS
	(1)	(2)	(3)	(4)
Share convictions t-1		-0.250**		0.033
		(0.101)		(0.021)
District harshness t-1	-0.055***		0.006**	
	(0.016)		(0.003)	
Observations	58745	58745	1103	1103
DepVarMean	1.357	1.357	3.417	3.417
F-stat excl. instrument		13.4		8.7
Clusters	225	225	268	268

Note: Relative to the baseline, the sample excludes district-year observations with a prior conviction rate equal to zero (i.e., no prior convictions among judges in that district-year). Reduced-form and 2SLS specifications otherwise follow the baseline (district and year fixed effects; survey regressions include marriage-duration fixed effects as well). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C A parametric model example

This appendix provides a simple parametric example consistent with the theoretical model in Section 5.2.

C.1 Functional forms

Escalation. Let escalation conditional on reporting be linear in severity:

$$\pi(s) = s, \qquad s \in [0, 1]. \tag{A1}$$

Reporting-cost heterogeneity. Wives differ in their private cost of reporting, κ (e.g. stigma, time, fear). Given severity s, a wife reports if the expected benefit from a successful case exceeds this cost:

$$\text{report} \iff q(s)\phi \geq \kappa, \qquad \text{where } q(s) = \mu\pi(s).$$

Hence the probability of reporting at severity s is

$$r(s) = \Pr(\kappa \leq q(s)\phi) = G(q(s)\phi) = G(\mu\pi(s)\phi)\,. \tag{A2}$$

To keep the example tractable while allowing flexible reporting elasticities, assume that in the relevant interior region the CDF takes the power form

$$G(x) = x^{\alpha} \qquad \text{for } x \in [0, 1], \qquad \alpha > 0. \tag{A3}$$

In the interior region where

$$x \equiv \mu\,\pi(s)\phi = \mu s\phi \in (0, 1), \tag{A4}$$

the reporting probability then becomes

$$r(s) = G(\mu\pi(s)\phi) = (\mu\pi(s)\phi)^{\alpha} = (\mu s\phi)^{\alpha}. \tag{A5}$$

The parameter α controls the curvature of reporting: $\alpha = 1$ implies linear reporting in x, while $\alpha > 1$ implies a more convex response: reporting is relatively low at small x but rises more quickly as x increases, see Figure A2.

C.2 Equilibrium severity

Under (A1) and (A5), the husband's objective (11) becomes

$$\begin{aligned} \max_{s \geq 0}\; bs - \frac{c}{2}s^2 - F\,\mu\,\pi(s)\,G(\mu\pi(s)\phi) &= \max_{s \geq 0}\; bs - \frac{c}{2}s^2 - F\,\mu\,s\,(\mu s\phi)^{\alpha} \\ &= \max_{s \geq 0}\; bs - \frac{c}{2}s^2 - F\phi^{\alpha}\,\mu^{\alpha+1}\,s^{\alpha+1}. \end{aligned} \tag{A6}$$

Figure A2: CDF with different α

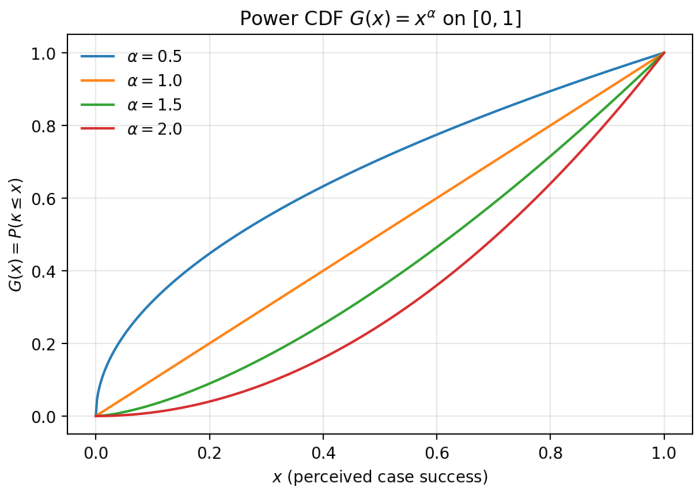

An interior optimum $s(\mu)$ satisfies the first-order condition

$$b - c\,s(\mu) - (\alpha + 1)\,F\phi^\alpha\,\mu^{\alpha+1}\,(s(\mu))^\alpha = 0. \tag{A7}$$

The last term is increasing in μ for any $s > 0$, implying $ds(\mu)/d\mu < 0$ under standard regularity conditions. When $\alpha = 1$, (A7) simplifies to a closed-form solution:

$$s(\mu) = \frac{b}{c + 2F\phi\mu^2}, \tag{A8}$$

which makes deterrence immediate since $ds(\mu)/d\mu < 0$ for $\mu > 0$.

C.3 Recorded crimes

Recorded crimes are defined as in the main text:

$$C(\mu) \equiv \pi(s(\mu))\,r(s(\mu)). \tag{A9}$$

Using (A1) and (A5), recorded crimes take the closed form

$$C(\mu) = s(\mu)\,(\mu\phi\,s(\mu))^\alpha = \mu^\alpha\phi^\alpha\,(s(\mu))^{\alpha+1}. \tag{A10}$$

Sign and when it can be positive. To express effects per one–percentage-point change in conviction probability (as in the empirical log specification), define $z \equiv 100\mu$ so that $dz = 1$ corresponds to a 1pp increase. Taking logs of (A10) and differentiating yields

$$\frac{d\log C(\mu)}{dz} = \frac{\alpha}{100\,\mu} + \frac{\alpha + 1}{100}\frac{d\log s(\mu)}{d\mu}. \tag{A11}$$

The first term is the *reporting/case-success* channel: holding severity fixed, higher μ raises the likelihood that a report leads to conviction, which increases reporting through $r(s) = (\mu s \phi)^{\alpha}$. The second term is the *deterrence* channel: higher μ reduces equilibrium severity ($ds/d\mu < 0$), which lowers both escalation $\pi(s)$ and reporting incentives that operate through severity.

Equation (A11) makes the ambiguity transparent. Recorded crimes increase with conviction probability at μ if and only if the reporting channel dominates deterrence:

$$\frac{\alpha}{\mu} > -(\alpha+1)\frac{d\log s\left(\mu\right)}{d\mu}. \tag{A12}$$

Thus, a positive response of recorded crimes is more likely when baseline conviction probabilities are low (making α/μ large) and when reporting is sufficiently elastic (larger α), and less likely when deterrence is strong (large $-\frac{d\log s}{d\mu}$).

C.4 Illustrative mapping to the empirical pattern

Finally we asks whether the parametric example can reproduce the qualitative empirical pattern: a negative survey-based effect on violence alongside a positive effect on recorded crimes.

Linking the survey semi-elasticity to the model. Using the estimates from the empirical analysis, we found survey-based semi-elasticity of violence onset $\widehat{\varepsilon}_V \approx -0.1507$ per one–percentage-point increase in the conviction rate. In the model with $\pi(s) = s$, the escalation/incident margin equals $\pi(s(\mu)) = s(\mu)$, so the model counterpart is

$$\left.\frac{d\log \pi(s\left(\mu\right))}{dz}\right|_{\mu=\mu_0} = \left.\frac{d\log s\left(\mu\right)}{dz}\right|_{\mu=\mu_0} \approx \widehat{\varepsilon}_V, \qquad z \equiv 100\mu.$$

Implied effect on recorded crimes. Evaluating (A11) at μ_0 and using $\frac{d\log s}{d\mu} = 100\,\frac{d\log s}{dz}$ gives the model-implied administrative semi-elasticity:

$$\left.\frac{d\log C(\mu)}{dz}\right|_{\mu=\mu_0} = \frac{\alpha}{100\,\mu_0} + (\alpha+1)\widehat{\varepsilon}_V. \tag{A13}$$

Thus, recorded crimes rise with conviction probability at μ_0 if and only if

$$\frac{\alpha}{100\,\mu_0} + (\alpha+1)\widehat{\varepsilon}_V > 0. \tag{A14}$$

Numerical implication at the sample mean. If the sample-mean conviction probability is $\mu_0 = 0.03732$ (i.e. 3.732%), then $1/(100\mu_0) \approx 0.2679$. With $\widehat{\varepsilon}_V \approx -0.1507$, condition (A14) becomes

$$0.2679\,\alpha - 0.1507(\alpha+1) > 0 \quad\Longleftrightarrow\quad \alpha > 1.29.$$

Hence, even given the estimated deterrence effect in the survey, the model can generate a positive effect on recorded crimes provided the reporting response is sufficiently elastic.

Finally, imposing the observed administrative semi-elasticity $\widehat{\beta}_{\text{admin}} = 0.0279$ in (A13) yields

$$0.0279 = \frac{\alpha}{3.732} + (\alpha + 1)(-0.1507) \quad \Rightarrow \quad \alpha \approx 1.52,$$

so a modestly convex reporting response (i.e. $\alpha > 1$) is sufficient to reconcile a negative survey response with a positive administrative response in this simple parameterization.[17]

[17] This mapping is local and illustrative rather than structural. The parameter α should be interpreted as a reduced-form summary of how reporting and recording respond to perceived case success; in practice, it may capture both heterogeneity in reporting costs and additional institutional frictions between reporting and administrative recording.

D Changes in social norms

We could also investigate changes in norms, as the NFHS asks questions in both the women's and men's modules about the acceptance of wife beating. Respondents are asked whether wife beating is justified if the wife: (i) goes out without telling her husband; (ii) neglects the children; (iii) argues with her husband; (iv) refuses sex; (v) burns the food; (vi) is unfaithful; or (vii) is disrespectful to her in-laws. Following Shreemoyee et al. (2025), we construct an indicator equal to one if a respondent justifies wife beating in *any* of the seven scenarios, and zero otherwise. This norms outcome is measured only at the time of the survey, and the NFHS does not elicit retrospective beliefs. As a result, we cannot construct a retrospective time series and instead observe (at most) two survey periods per district (2015–2016 and 2019–2021).

Given this constraint, we apply the baseline specification, effectively comparing respondents from two survey waves within each district. The small number of periods limits statistical power, so we treat the estimates as suggestive. For women, the reduced-form estimates in Column 1 of Table A8 indicate that greater judge harshness in the previous year causes lower acceptance of wife beating: a one–percentage-point increase in judge harshness reduces acceptance by 0.4 percentage points, or 0.9% of the sample mean ($p < 0.10$). For men, we find no evidence of changing norms.[18]

Table A8: Acceptance of wife beating

	Wives (1)	Husbands (2)
District harshness t-1	-0.004* (0.002)	-0.000 (0.003)
Observations	79080	61557
Dep.var.mean	0.468	0.381

Note: The dependent variable is an indicator equal to one if the respondent reports that wife beating is justified in *any* of the seven scenarios and zero otherwise. Column (1) uses the women's module (wives) and Column (2) uses the men's module (husbands). Both regressions follow the baseline reduced-form design with district and survey-year fixed effects. "Dep.var.mean" reports the sample mean of the dependent variable in each column. Standard errors are clustered at the district level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

[18] With only two survey waves, the instrument has insufficient within-district time variation, and we have no first stage. We therefore do not report 2SLS estimates.

E Extra tables

Table A9: Self-reported media use, 2SLS estimates

	Violence onset			Reported crimes		
	News-paper	Radio	TV	News-paper	Radio	TV
	(1)	(2)	(3)	(4)	(5)	(6)
Share convictions t-1	-0.224***	-0.218**	-0.223**	0.035	0.028*	0.039**
	(0.084)	(0.087)	(0.087)	(0.021)	(0.017)	(0.018)
Share convictions t-1	-0.013	-0.255	-0.006	-0.038	-0.003	-0.027
$\times$Media use	(0.159)	(0.229)	(0.119)	(0.031)	(0.016)	(0.019)
Observations	102614	102614	102614	1875	1875	1875
Dep.var.mean	1.486	1.486	1.486	3.324	3.324	3.324
Clusters	384	384	384	435	435	435
F-stat excl. instrument						
Share convictions t-1	8.9	9.6	10.8	6.6	8.3	6.7
Interaction	4.3	1.0	4.0	3.2	2.5	3.7
Media, level	district	district	district	district	district	district

Note: The table tests for heterogeneity by local media exposure at the level of districts. Each column reports a 2SLS regression where we instrument for both the conviction rate and its interaction with augmented the media measure, as indicated by the column headings. The specification otherwise follows the baseline specification (district and year fixed effects; survey regressions also include marriage-duration fixed effects). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A10: Access to community radio states and violence, 2SLS estimates

	Violence onset (1)	**Reported crimes** (2)
Share convictions t-1	-0.220**	0.023*
	(0.089)	(0.014)
Radio	-0.142	0.039**
	(0.129)	(0.019)
Share convictions t-1	-0.009	-0.004
$\times$Radio	(0.022)	(0.003)
Observations	102614	1623
Dep.var.mean	1.486	3.343
Clusters	384	370
F-stat excl. instrument		
Share convictions t-1	11.9	12.0
Interaction	3.7	6.3
Radio, level	clusters	district

Note: The table tests for heterogeneity by community radio availability at the level of survey clusters/districts. Each column reports a 2SLS regression where we instrument for both the conviction rate and its interaction with augmented the media measure. The specification otherwise follows the baseline specification (district and year fixed effects; survey regressions also include marriage-duration fixed effects). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A11: Access to local newspapers and violence, 2SLS estimates

	Violence onset (1)	**Reported crimes** (2)
Share convictions t-1	-0.283***	0.021
	(0.090)	(0.013)
Share convictions t-1	-0.088	0.001
$\times$ Newspaper circulation p.c. (std)	(0.077)	(0.005)
Observations	94127	1647
Mean.dep.var	1.409	3.338
Clusters	380	377
F-stat excl. instrument		
Share convictions t-1	9.1	11.3
Interaction	15.0	39.0
Newspaper, level	district	district

Note: The table test for heterogeneity by local newspaper circulation at the level of districts. Each column reports a 2SLS regression where we instrument for both the conviction rate and its interaction with augmented the newspaper measure. The specification otherwise follows the baseline specification (district and year fixed effects; survey regressions also include marriage-duration fixed effects). Standard errors are clustered by district. "Dep.var.mean" reports the mean of the dependent variable.
Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.