Conflict and the Persistence of Ethnic Bias*

Moses Shayo and Asaf Zussman

The Hebrew University of Jerusalem

May 9, 2016

Abstract

How persistent are the effects of conflict on bias towards co-ethnics? What are the channels of persistence? We employ a measure of ethnic bias derived from decisions made by Israeli Arab and Jewish judges to study the levels and determinants of bias during the 2000-04 conflict and its aftermath (2007-10). Despite the fall in violence, we find no evidence of a general attenuation in bias. Furthermore, bias remains positively associated with past intensity of violence in different localities. This persistence does not appear to be due to judges' personal exposure to violence but rather to different dynamics in afflicted areas.

* We thank Kate Antonovics, Eli Berman, Daniel Chen, Gordon Dahl, Ernesto Dal Bó, Badi Hasisi, Maria Petrova, Dan Posner, Nicholas Sambanis and participants of seminars at Ben Gurion University, ETH Zurich, Hebrew University, Max Planck Institute Bonn, NBER Summer Institute 2014, the Political Economy of Conflicts and Development conference, Stanford, UCLA, UCSD and the Vienna Workshop on Behavioral Public Economics for many valuable comments and suggestions. Excellent research assistance was provided by Raya Adani, Eli Bing, Ana Danieli, Eitan Gadon, Noam Goldman, Lee Goren, Alon Grembeck, Vered Isman, Ohad Katz, Noa Litmanovitz, Netta Porzycki, Niva Porzycki, Lilach Rapaport, Ittai Shacham, Michal Shamir, Elad Shapira, Karin Telio and Rozi Tshuva. Financial support for the project was generously provided by the Israel Science Foundation (grant 1411/12), the I-CORE Program of the Planning and Budgeting Committee at the ISF (grant no. 1821/12) and the Maurice Falk Institute. Shayo gratefully acknowledges support from the European Union, ERC Starting Grant Agreement no. 336659.

1. Introduction

This paper seeks to contribute to our understanding of the legacies of ethnic conflict. Our focus is on ingroup bias, one of the key behavioral manifestations of social identification. While the fluidity of social identities has been well documented both in political science and in social psychology, when it comes to large scale ethnic violence a widespread view is that conflict tends to harden ethnic identities – perhaps irreversibly.[1] Previous work established a strong short-term effect of ethnic violence in a locality on the extent of ethnic bias in that locality (Shayo and Zussman 2011). This paper examines whether the effects of conflict persist even when violence subsides or whether, on the contrary, ethnic bias in previously afflicted areas converges towards the lower levels of bias observed in non-violent localities. Furthermore, we seek to better understand the mechanisms that might cause the effects of conflict to persist. One possibility, consistent with the view that conflict irreversibly solidifies ethnic identities, is that the personal experience of ethnic violence has an enduring effect on individuals' social preferences. A second possibility is that individual identities are in fact quite fluid, and yet *in equilibrium* identities can be very stable. This may happen because individuals coordinate on patterns of identification and inter-group relations (see e.g. Laitin 2007) or because

[1] Chandra (2012, ch.1) provides an overview of "constructivist" accounts of ethnic identity in political science and Shayo (2009) reviews the experimental literature on the endogeneity of social identities. The summer 2004 volume of *Security Studies* contains a collection of essays debating Kaufmann's (1996) thesis that "in ethnic wars both hyper-nationalist mobilization rhetoric and real atrocities harden ethnic identities to the point that cross-ethnic political appeals are unlikely to be made and even less likely to be heard" (p. 137) and his argument for partition as the solution to ethnic conflicts.

conflictual relationships and ethnic identification tend to reinforce each other (see e.g. Sambanis and Shayo 2013).

We define ethnic bias as the preferential treatment awarded to members of one's ethnic group. A measure of ethnic bias should arguably have the following two properties. First, it would be based on behavior in a natural setting. Second, it would reliably capture bias rather than unobserved factors. Consider the first property. No doubt, survey data as well as carefully constructed lab and lab-in-the-field experiments have yielded crucial insights into the determinants of ethnic bias. Nonetheless, survey responses are sometimes hard to interpret and may have little predictive power with respect to behavior.[2] Experiments have the advantage of cleanly identifying biases in behavior and exploring their possible causes. Nonetheless, external validity may sometimes be a concern, e.g., due to the participants being nonprofessional decision makers, acting in an artificial context, possibly characterized by low stakes.

Turning to the second property, a large literature seeks to measure ethnic, racial, and gender biases based on behavior in natural settings (e.g. courts and the labor market). However, a serious challenge facing this literature is that outcomes that seem to reflect bias may actually be driven by other factors which are correlated with ethnicity, race or gender. For instance, wage disparities along ethnic lines could be due to unobserved variables (such as ability, motivation, and quality of education) rather than to ethnically-biased employers. A similar problem plagues empirical legal studies, where many case

[2] See Zussman (2013) and references therein on the relation between stated attitudes and discriminatory behavior.

characteristics that might affect judicial decisions are unobserved by the researcher. For recent reviews, see Bertrand and Duflo (forthcoming) and Neumark (forthcoming).

Here, we utilize the measure of ethnic bias developed by Shayo and Zussman (2011) which has these two properties. The measure is derived from rulings made by Arab and Jewish judges in Israeli small claims courts in cases involving either an Arab plaintiff suing a Jewish defendant or a Jewish plaintiff suing an Arab defendant. In these courts, cases are effectively randomly assigned to judges. This allows us to estimate the extent of judicial ethnic bias – which is probably a lower bound on the level of ethnic bias in the wider society. Data on fatalities from the Israeli-Palestinian conflict serve as a proxy for exogenous temporal and spatial variation in exposure to ethnic violence. This allows us to relate the extent of the bias to the court's and the judge's exposure to violence, while controlling for factors operating at the national level that may affect both judicial decisions and conflict intensity.

Figure 1 displays the number of Israeli civilian fatalities from the Israeli-Palestinian conflict during 1998–2010.[3] Looking at the figure, one can distinguish between three periods. The period from 1998 to late 2000 was relatively peaceful. The period from September 2000 to the end of 2004 is known as the Second Palestinian Intifada (uprising). This period was characterized by intense conflict between Israel and the Palestinians in the West Bank and the Gaza Strip, with much of the violence taking place within Israel. Violence subsided in

[3] We focus on fatalities within Israel proper (i.e. excluding fatalities in the Occupied Territories and Lebanon) since this is where the courts we study are located. Effects of violence in other areas cannot be identified and will be absorbed in our empirical analyses by time fixed effects and overall trends.

subsequent years, with 2007-2010 characterized again by relative calm. Henceforth we refer to 2000-2004 as the "conflict period" and to 2007-2010 as the "post-conflict period" but it is worth emphasizing that the Israeli-Palestinian conflict, in the wider sense of the term, did not start in 2000 and has not reached a peaceful conclusion yet. In particular, the post-Intifada years saw an increase in the severity of the rocket threat facing Israel. While rockets launched from the Gaza Strip caused few civilian fatalities inside Israel, the rocket threat could still have affected social and political attitudes. Indeed, Getmansky and Zeitzoff (2014) find that Israelis residing within the rocket range are more likely to vote for right-wing parties.

[Figure 1]

Shayo and Zussman (2011) show that during the conflict period judges exhibited a significant degree of ethnic bias. A case was around nineteen percentage points more likely to be accepted when assigned to a judge of the same ethnicity as the plaintiff. Furthermore, the extent of the bias was strongly associated with the intensity of ethnic violence in the vicinity of the court in the period immediately preceding the trial. Courts that had seen limited violence exhibited little bias. Based on these findings, one might conjecture that a decline in violence should lead to a reduction in bias, especially in those courts where violence dropped most significantly. This paper examines the actual legacy of the conflict.

We make three main contributions. The first is to provide an answer to a simple factual question: can a sharp decline in violence indeed reduce ethnic bias? This is a fundamental question for any theory of ethnic conflict and reconciliation. Our answer is sobering. We find no evidence of an overall attenuation in bias as violence subsides – even among an educated elite operating under a strong norm of equality before the law. In the post-conflict period, claims were still around eighteen percentage points more likely to be accepted when assigned to a judge of the same ethnicity as the plaintiff. Furthermore, we observe no downward trend in ethnic bias within the post-conflict period. While the 2007-2010 period clearly wasn't the most peaceful of times, and other events could have fueled ethnic bias, the drop in ethnic violence within Israel documented in Figure 1 is quite dramatic. The fact that it was not associated with a reduction in bias is important and informative in itself.

The second contribution concerns the persistence of the effects of exposure to violence. Our evidence indicates that ethnic bias in the post-conflict period remains positively associated with local exposure to violence during the height of the conflict. This association is not due to any other observable differences across the localities where the courts are located – despite the fact that some such differences are associated with the level of bias (courts located in larger towns exhibit less bias and there is some indication that so are courts located in mixed Arab-Jewish towns). Moreover, placebo tests show that the intensity of violence itself is not associated with pre-existing levels of ethnic bias in the courts. However, while the association between judicial bias and past violence persists, it declines in strength compared to the conflict period. Taken together, this pattern and the first set of results suggest some convergence across courts toward a relatively high level of judicial bias in the post-conflict years.

The third, and perhaps the most significant contribution concerns the sources of persistence. Broadly, ethnic violence could have a persistent effect on the legal system via at least two distinct channels. The first operates through the personal experience and memories of individual judges. This channel is consistent with studies indicating that exposure to political violence produces psychological distress and increased threat perceptions, which may lead to increased militancy and to hostility towards minority groups (Canetti-Nisim et al. 2009, Hirsch-Hoefler et al. forthcoming). Such effects could conceivably persist even when the level of threat has subsided. Thus, judges who were exposed to violence are, in a sense, “scarred”. But violence might produce enduring changes in exposed courts independently of the individual experiences of present-day judges. Violence during the time of the conflict can change the patterns of inter-ethnic relations in the region where the court is situated. Increased tensions between the groups raise the salience of ethnic cleavages in the region, thereby increasing the likelihood of ethnic identification which in turn reinforces inter-ethnic tensions (see Sambanis and Shayo 2013 for a theoretical analysis of this process). Thus, even individuals who were not personally exposed to violence during the conflict may exhibit ethnic bias. To shed light on the sources of persistence, we collect data on judges’ personal history of exposure to violence. This allows us to disentangle the two potential mechanisms. The evidence seems more in line with the second: the persistent effect of exposure to violence that we identify operates primarily at the local (court) level rather than at the individual (judge) level.[4]

[4] Whether the individual or the social mechanism is at work can be important to help guide post-conflict policies. Of particular relevance are Community Driven Reconstruction programs which seek to enhance local cooperation by encouraging individuals in a community to work together; and Truth and Reconciliation programs which emphasize psychological healing (see, respectively, Fearon, Humphreys and

The paper relates to a recent line of research that has begun to uncover the effects of conflict on political attitudes (Berrebi and Klor 2008; Gould and Klor 2010; Lyall et al 2013; Getmansky and Zeitzoff 2014), social, risk, and time preferences (Voors et al. 2012; Cassar, Grosjean and Whitt 2013; Rohner, Thoenig and Zilibotti 2013; Callen et al. 2014; Gilligan, Pasquale and Samii 2014) and political participation (Bellows and Miguel 2009; Blattman 2009). Our contribution to this literature is in studying the lasting impact of ethnic conflict on intergroup relations and on the functioning of institutions. While some of the existing literature finds that exposure to violence can increase pro-social behavior or public spiritedness, it is unclear whether such favorable changes in behavior are biased towards one's own ethnic group. Ingroup bias plays a crucial role in intergroup relations and may contribute to the emergence of vicious cycles where violence breeds ethnic identification and ethnic identification intensifies conflict.

Concerning the relation between conflict and institutions, what we know is largely based on cross-country comparisons. Collier et al. (2003) report that indices of democratic institutions (from Polity IV) and of political freedom (from Freedom House) tend to deteriorate in the decade following civil war. Besley and Persson (2008) find that tax collection as a percentage of GDP (a proxy for state capacity) is lower in countries that experienced internal conflict. While instructive, such results are hard to interpret: countries with bad or deteriorating institutions may be more likely to experience war for many reasons. To our knowledge, this paper is the first to employ micro data to analyze the enduring effects of ethnic violence on the workings of legal institutions.

Weinstein (forthcoming) and Cilliers, Dube and Siddiqi (2015) for recent discussions and program evaluations).

Several recent studies examine the persistence of biases, norms and preferences. The evidence suggests, for example, that trust, anti-Semitism, political attitudes and perceptions of gender roles can persist over very long horizons (see, respectively, Nunn and Wantchekon 2011; Voigtländer and Voth 2012; Acharya, Blackwell and Sen 2015; Alesina and Fuchs-Schündeln 2007; Alesina, Giuliano and Nunn 2013). Besley and Reynal-Querol (2014) find that the intensity of conflict in pre-colonial Africa is positively correlated with present-day ethnic (as opposed to national) identification and negatively correlated with inter-group trust. But while norms and identities can persist, they can and do change, sometimes quite rapidly. Here, we closely examine the extent to which one crucial form of bias persists in the face of a rather dramatic change in the environment, and study the underlying mechanism.

In terms of methodology, the paper builds on the (largely experimental) literature on social identity and ingroup bias (e.g. Tajfel et al. 1971; Bornstein 2003; Habyarimana et al. 2007; Chen and Li 2009; Fong and Luttmer 2009; Klor and Shayo 2010). One of the workhorse tools for studying ingroup bias has been the Minimal Group Paradigm introduced by Tajfel et al. (1971). In these lab experiments, subjects are assigned to groups and are then asked to make allocation decisions between two other anonymous subjects, one from their group (the ingroup) and one from the outgroup. Typically, subjects exhibit ingroup bias. Our setting is unique in that it replicates some of the key properties of these experiments, but with professional decision makers acting in their natural environment and with ethnic (rather than minimal) groups.

Finally, the study relates to the literature on ethnic and racial bias in the legal system (e.g. Glaeser and Sacerdote 2003; Abrams, Bertrand, and

Mullainathan 2012; Anwar, Bayer, and Hjalmarsson 2012; Alesina and La Ferrara 2014, Grossman et al. 2016). Beyond documenting persistent ethnic bias, our contribution to this literature is in helping to understand how such biases are shaped by the history of intergroup relations.

2. Institutional and historical background

We analyze judicial decisions involving Arabs and Jews in Israeli small claims courts in 2000–2010. During this time, Arab citizens account for around 20% of Israel's population while Jewish citizens account for 76%.[5] Arab-Jewish relations within Israel are closely linked to developments in relations between Israel and the Palestinians in the Occupied Territories. Arab Israelis and the Palestinians of the Occupied Territories share a common heritage and many in the latter group have roots in towns and villages located in present-day Israel.

Small claims courts handle civil cases between private litigants, such as those involving minor traffic accidents. The courts have a cap on the amount of money they can award. This cap increased from NIS 17,800 in 2000 to NIS 31,200 in 2010 (roughly $US 4,000-8,000). The judge can also award legal expenses (this is entirely discretionary). The legal procedure is relatively simple. It starts when the plaintiff files a claim at the court, provides supporting documentation, and pays a small fee. The defendant is then notified and instructed to provide a defense statement within fifteen days. After a claim has been filed, the case is assigned a trial date and a judge. Due to a backlog in the system, trials are scheduled several months in advance. Importantly, each case is assigned to the

[5] Israeli Central Bureau of Statistics, *Statistical Abstract of Israel 2013*. The remaining 4% are mostly immigrants from the former Soviet Union who are not officially classified as Jews.

first available slot. This means that the assignment of judges to cases within a given court is in principle orthogonal to characteristics of the case (we test this below).

The judge receives the case materials no earlier than a week before the trial. The trial itself typically lasts only a few minutes and the judge issues a ruling within seven days. Importantly, the litigants (rather than lawyers) present their case.

This setting has several attractive features for measuring ethnic bias. Most importantly, in these courts an effectively randomly assigned decision maker needs to make an allocation decision between two individuals, one of whom belongs to her ethnic group while the other does not. Compared to lab experiments designed to measure ingroup bias (e.g., Tajfel et al. 1971; Chen and Li 2009), decision makers are professional, the monetary stakes are high, and the groups to which the agents belong (Arabs and Jews) are not artificially constructed in the lab. Since decisions have to be made within a week of the trial, we also know their approximate timing. This allows us to relate the decisions to the context in which they are made and in particular to the intensity of the conflict.

We concentrate on two periods. The first is the conflict known as the Second Intifada, which started in September 2000 and lasted until the end of 2004.[6] It was characterized by intense violence between Israelis and Palestinians

[6] There is no agreed-upon date for the end of the Intifada. Two prominent candidate dates are November 11, 2004, when Palestinian leader Yasser Arafat died, and February 8, 2005, when the Sharm el-Sheikh Summit was convened to end the Intifada. We use December 31, 2004, as the cutoff date for the end of the Intifada (none of the results depend on this decision).

and claimed the lives of thousands. A major part of Palestinian violence took the form of suicide bombings carried out within Israel proper, where the courts we study are located. Importantly, this violence was not generated by the local population. In particular, participation of Arab Israelis in terrorist attacks was negligible in scale throughout the conflict period. Violence reached a peak in 2002 and then gradually declined (see Figure 1). The second period we study (2007–2010) was relatively tranquil and saw only 36 civilian fatalities within Israel compared to 514 in the first period.[7]

3. Data

This section briefly outlines the data collection procedures. Details are provided in Appendix A which also reports comprehensive descriptive statistics.

Judicial decisions.

The main source of data is online transcripts of judicial decisions (rulings). We cover the universe of available documents (26,444 from 2000–2004 and 28,576 from 2007–2010). For each document we code the ethnicity of private litigants based on their names and the relative frequency of Arab and Jewish names in the Israel Population Registry. We keep only "mixed cases": those where at least one private plaintiff and one private defendant are of different ethnicities (N=4,038).[8] For these cases we conduct a comprehensive analysis of the

[7] The post-Intifada period did see some outbreaks of violence, including the Second Lebanon War in the summer of 2006, which did not involve the Palestinians; and the winter of 2008–2009 Israeli military operation in the Gaza Strip. Both events involved a relatively small number of casualties within Israel. In the analysis that follows we control for these and other incidents of violence in the post-Intifada period.

[8] Focusing on mixed cases allows us to reliably measure ethnic bias, which is our primary interest. A comprehensive analysis of the universe of cases could perhaps allow us to answer other questions but would be prohibitively costly.

documents, coding information on the court, judge, litigants (both plaintiffs and defendants), claim characteristics and judicial outcomes. We obtain additional socio-demographic information on judges from their biographies.

The main analysis in this paper excludes cases that were settled outside the court (325 cases) or withdrawn (303) as well as cases that have multiple plaintiffs (or defendants) such that one plaintiff (or defendant) is Jewish and another is Arab (305).[9] This leaves us with 3,153 cases, 1,748 for 2000–2004 and 1,405 for 2007–2010.

Table 1 shows for each period the percentage of cases by the ethnicity of the judge, plaintiff, and defendant. Overall, 29% of the cases in our data were ruled by Arab judges. This share dropped from 31% to 25% across the two periods. There is also a small decline in the share of Arab plaintiffs, from 44% in the conflict period to 41% in the post-conflict period.

[Table 1]

Our main measure of trial outcome is an indicator for whether the claim was accepted. Out of the 3,153 cases in our main sample, 2,300 (73%) are coded as accepted. In Appendix D we also provide results using alternative outcome measures such as monetary compensation and legal expenses.

[9] We find no evidence for differences in the likelihood that plaintiffs settle cases outside the court or withdraw them when assigned a judge of the same or opposite ethnicity. We similarly find no evidence that such differences are associated with court exposure to violence during the conflict. See Appendix B.

Exposure to violence.

We employ data on all Palestinian politically motivated fatal attacks inside Israel. For each attack we have information about date, location, and number of civilian fatalities.

Our first set of measures of exposure to violence is at the level of the court. These measures are based on the number of fatalities from attacks that occurred in the vicinity of the court during the conflict period. Vicinity is defined by three alternative geographical units: *natural area* (smallest geographical unit around the court) *sub-district* and *district*. Our data spans 24 natural areas, 15 sub-districts, and 6 districts.

The second set of measures is at the level of the judge. We focus on exposure at the judge's place of employment. We use several sources to compile information on judges' employment history since 2000 (both as judges and in other jobs such as lawyers). Our procedure yields monthly location data for the entire Intifada period for 196 (82%) of the 240 judges in our sample. For an additional 37 judges we have partial information and for 7 judges we have no location information for the conflict period. Merging the location information with the fatalities data provides a measure of the number of fatalities each judge was exposed to in her place of employment in each month of the conflict.

4. Ethnic bias across periods

Are claims more likely to be accepted when assigned to a judge of the same ethnicity as the plaintiff? The random nature of the assignment of cases to judges within a court provides a straightforward method for credibly addressing this question. Our identification assumption is that given the court, the ethnicity of the plaintiff, and the ethnicity of the judge, cases assigned to a judge of the same ethnicity as the plaintiff are not systematically different from cases assigned to a judge of a different ethnicity. Notice in particular that we allow plaintiffs from different ethnicities to file cases with different characteristics. We also allow judges of different ethnicities to receive cases with different characteristics. Appendix C evaluates the validity of our identification assumption using the observed case characteristics (e.g. the number and characteristics of plaintiffs and defendants, the subject of the claim and other case characteristics). The evidence strongly supports the validity of the identification assumption.

4.1 Results

We first present the raw data and then proceed to an econometric analysis. Figure 2 displays the share of claims accepted for each period by judge and plaintiff ethnicity. The left box replicates Shayo and Zussman's (2011) results from the conflict period and serves to illustrate the methodology. The left pair of bars in this box pertains to cases where the plaintiff is Jewish and the defendant is Arab. Jewish judges accept 79.1% of these cases while Arab judges accept 71.7%. In itself, this difference of 7.4 percentage points is not necessarily evidence of ethnic bias: for example, it may be the case that Jewish judges are more inclined to accept claims than Arab judges. However, if this were the only reason for the difference, we would expect to observe a similar pattern regardless of plaintiff ethnicity. In fact, the right pair of bars in the same box shows that when the plaintiff is Arab, the pattern is reversed: Jewish judges are 10.6 percentage points *less* likely than Arab judges to accept such claims. The difference in these differences – 18 percentage points – provides an indication of the extent of ethnic bias (i.e., by how much Jewish judges are more likely than their Arab colleagues to accept a claim filed by a Jewish plaintiff rather than by an Arab one). From now on we will use this measure to track the evolution and identify sources of ethnic bias. It should be emphasized that this is a measure of *overall* bias: absent an ethnicity-free benchmark, it is impossible to say whether and to what extent Jewish judges favor Jewish litigants and Arab judges favor Arab litigants.

[Figure 2]

We can now turn to the post-conflict period (right box). Compared to the conflict period, we do observe some changes (e.g., Jewish judges accept a lower share of claims, while Arab judges accept a higher share). Remarkably, however, the overall level of bias is stable at 17.3 percentage points. Note, again, that it is impossible to say to what extent the bias is due to the behavior of Arab or Jewish judges. For example, looking at the right box, one might think that the bias is mainly driven by the behavior of Jewish judges (who accept a significantly lower proportion of cases filed by Arab plaintiffs). However, one cannot rule out the possibility that cases filed by Jewish plaintiffs are systematically stronger than those filed by Arab plaintiffs, in which case the decisions by Arab judges are biased.

A more detailed look at the evolution of bias over time is provided in Figure 3. The figure presents our difference-in-differences measure of ethnic bias, but this time it is calculated in two-year moving windows centered around a particular date. We start with the first observation in each period (conflict and post-conflict) and advance in 30-day intervals until the end of the period. To reduce noise, the analysis is restricted to windows with at least 100 observations. The results show the bias fluctuating around 15–20 percentage points during both the conflict period and the post-conflict period, with no clear downward trend following the decline in violence.

[Figure 3]

We now turn to an econometric investigation. We start with a benchmark difference-in-differences specification:

$$(1) \qquad y_{ijct} = \alpha_0 + \alpha_1 ArabPlaintiff_i + \alpha_2 ArabJudge_i$$

$$+\alpha_3 ArabPlaintiff * ArabJudge_i + \delta_c + \varepsilon_{ijct}$$

where y_{ijct} is the outcome of case i assigned to judge j, in court c, at time t; *ArabPlaintiff*, *ArabJudge* and the interaction term *ArabPlaintiff*∗*ArabJudge* are indicator variables; δ_c is a court fixed effect; and ε_{ijct} is an error term clustered within judge.

Consistent with our identification assumption, equation (1) allows for two possible differences across ethnic groups that do not necessarily indicate ethnic bias. First, it is possible that claims submitted by Arab plaintiffs have different unobserved characteristics than those submitted by Jewish plaintiffs. Thus, α_1 may be nonzero even in the absence of ethnic bias. Second, it is possible that Arab and Jewish judges are differently inclined to accept claims. In other words, α_2 may be nonzero even in the absence of ethnic bias. Judicial ethnic bias is captured by α_3. This coefficient is similar to the difference-in-differences measure used in Figures 2 and 3, except that we include court fixed effects to allow for differences in acceptance rates across courts. Columns 1–2 in Table 2 present these benchmark results for the conflict period and the post-conflict period, respectively. Consistent with Figure 2, the estimated ethnic bias is similar in magnitude across periods (and highly statistically significant in both).

[Table 2]

In columns 3–4 we augment equation (1) with a large set of controls. We now estimate:

$$(2) \qquad y_{ijct} = \alpha_0 + \alpha_1 ArabPlaintiff_i + \alpha_3 ArabPlaintiff * ArabJudge_i + \delta_c$$

$$+\gamma_j + \lambda tenure_{jt} + X_i'\beta + \eta_t + \varepsilon_{ijct}$$

where γ_j is a judge fixed effect and $tenure_{jt}$ is the judge's tenure at the job.[10] The vector X_i is a list of case-specific controls (see notes to Table 2) and η_t is a vector of year, month, and day-of-week indicators. Compared to columns 1–2, the explanatory power of the regressions increases substantially. However, consistent with random assignment of cases to judges, the estimates of ethnic bias hardly change.[11]

While the post-conflict period saw a dramatic drop in civilian fatalities within Israel, it also saw increase in the severity of the rocket threat from the Gaza Strip facing southern Israel. This contemporaneous threat could have affected judicial decisions (cf. Getmansky and Zeitzoff 2014). However, excluding

[10] Adding the judge fixed effect implies dropping the $ArabJudge$ indicator from the model since the fixed effect picks up any time-invariant judge characteristic. We keep the court fixed effect δ_c as some judges move between courts.

[11] The results in Table 2 are not driven by the decisions of a single judge. To check this, we repeat the analysis of columns 3 and 4, each time excluding decisions made by a single judge. Estimates of ethnic bias range between 0.158 and 0.207 in the conflict period and between 0.156 and 0.202 in the post-conflict period.

all cases from courts located within rocket range in the south during this period has little effect on the estimated bias (point estimates are 0.179 and 0.180 in the specifications of columns 2 and 4).[12] We can similarly exclude all courts that were within Hezbullah's rocket range during the Second Lebanon War in 2006. This leads to a sharp reduction in the number of observations in the post-conflict period. Despite the very limited size of the remaining sample (N=194), the estimated bias remains statistically significant and if anything increases in size (point estimates are 0.399 and 0.442 in the specifications of columns 2 and 4).

In Appendix D we estimate equation (2) using four alternative outcome measures such as the monetary compensation awarded by the judge. The results are in line with those of Table 2: for the most part, the estimated bias is similar across periods. The main exception is that the bias measured in net monetary compensation decreases from NIS 926 (about $210) during the conflict period to 391 during the post-conflict period (panel A, columns 3–4). However, this decrease is driven by a small number of cases involving extraordinarily high monetary compensation. Excluding from the analysis the top and bottom one percent of cases in terms of the outcome variable, we find no evidence of a decline in bias across periods (panel B).

Before turning to an analysis of the relationship between judicial bias and violence, several issues are worth discussing. First, given that Arab and Jewish judges have different personal characteristics, it is possible that the estimated α_3 in equation (2) reflects these differences rather than ethnic bias per se. For example, Arab judges are more likely than Jewish judges to be male (66% vs.

[12] The potentially affected courts are located in Ashdod, Ashqelon, Be'er Sheva and Qiryat Gat. Removing them reduces the sample size to 1391.

49%). If for some reason male judges systematically treat Arab plaintiffs differently from Jewish plaintiffs, then α_3 would pick up this difference and hence would not be an accurate measure of ethnic bias. To address this possibility, in Appendix E we augment equation (2) with interactions between the $ArabPlaintiff$ indicator and judge characteristics. None of these added interactions meaningfully affect our estimate of ethnic bias.

Second, one could argue that the empirical pattern we observe is not due to preferential treatment of members of one's own group but rather to difficulties in communication between judge and plaintiff when they do not share the same mother tongue. Even if true, our results point to a severe malfunction in the legal system, as similar cases are not treated alike. However, it is hard to see why technical differences in the quality of communication between Arabs and Jews should respond (as we show below) to the intensity of violence in the vicinity of the court. Note also that cases are typically very simple (e.g., “fender-bender” accidents) and the decision is basically about whose version of the events to accept and as such should not rely on linguistic subtleties.

A third issue relates to the possibility that the strength of claims submitted by plaintiffs might depend on the ethnic makeup of the court. For example, Arab plaintiffs may file more “marginal” claims – and Jewish plaintiffs only “strong” claims – the higher is the share of Arab judges in the court. This can generate downward bias in our estimate of ethnic bias. To help assess this possibility we collect data on the ethnicity of all judges in each court and year, including judges that did not end up ruling in mixed-ethnicity cases. The data are derived from the official biographies of all current and retired judges and the computerized archive of judicial decisions described in Appendix A. We then augment equation (2) with two variables: the share of Arab judges in court c in year t and the interaction of this share with the $ArabPlaintiff$ indicator. If the above hypothesis is correct, we would expect the likelihood of a claim being accepted to be positively associated with the share variable and negatively associated with the interaction term. The two coefficients turn out to be statistically insignificant while the estimate of ethnic bias rises slightly.[13]

Finally, Appendix F examines whether the degree of ethnic bias varies with judge characteristics (e.g. do more experienced or more educated judges exhibit less bias). Overall, we find little evidence for systematic differences, although there is some indication that within the post-conflict period older and male judges are less biased.

5. Judicial bias and court exposure to violence

So far we have established that the dramatic drop in fatalities in 2007-2010 was not accompanied by a decline in ethnic bias among judges. However, bias could be affected by different country-wide factors that vary over time. In this section we therefore explore whether judicial bias in the post-conflict period is associated with a court's *local* exposure to violence during the conflict. This allows us to better identify the causal effect of violence, as we can keep constant any country-wide determinants of bias. Note however that this approach probably yields an underestimate of the overall effect of violence, since Israel is

[13] Recall also that we find no evidence of selective withdrawal of cases or their settlement outside of court depending on the ethnicity of the judge – which could similarly bias our estimate of ethnic bias (Appendix B).

a small country where national media outlets prominently cover any act of political violence.

Our identification assumption is that as local exposure to violence during the conflict period increases, cases assigned to a judge of the same ethnicity as the plaintiff do not become systematically different from cases assigned to a judge of a different ethnicity. Appendix G examines and finds support for the validity of this assumption.

5.1 Results

Figure 4 illustrates some broad patterns in the data. Using the same methodology as Figure 3, it shows simple (nonparametric) difference-in-differences estimates of judicial bias. We split the sample into two groups of courts based on the court's exposure to violence during the conflict. The first group ("low violence courts") consists of eight courts that saw no fatalities in their respective natural areas during the entire conflict period. The second group consists of all other courts.[14] During the conflict period, courts that saw no violence are consistently less biased than the other courts. In fact, the overall level of bias in this period is 0.05 in the former group and 0.23 in the latter. After the violence subsides, the difference between the two groups appears to narrow. The overall post-conflict bias is 0.15 in the low-violence courts and 0.23

[14] We have 915 cases in the low-violence group of courts (454 of them in the conflict period) and 2,238 in the high-violence group (1,294 of them in the conflict period). The distinction between the two groups is for illustrative purposes only. Some of the point estimates for the high-violence courts are based on samples that include cases in courts that are yet to experience fatalities. Nonetheless, the figure is useful as it keeps the set of courts in each group fixed. In Appendix H (discussed below) we conduct placebo tests to check whether courts exposed to more violence were more biased to begin with. We find no evidence of this being the case.

in the high-violence courts. Contrary to what one might have expected, the two groups seem to converge to the level of bias previously associated with the *high-violence* courts.

[Figure 4]

We now examine more closely the relationship between exposure to violence and bias in the two periods. Consider the following equation:

$$(3)\ y_{ijct} = \alpha_0 + \alpha_1 ArabPlaintiff_i + \alpha_3 ArabPlaintiff * ArabJudge_i + \theta_0 Exposure_{ct} + \theta_1 Exposure_{ct} * ArabPlaintiff_i + \theta_2 Exposure_{ct} * ArabJudge_i + \theta_3 Exposure_{ct} * ArabPlaintiff * ArabJudge_i + \delta_c + \gamma_j + \lambda tenure_{jt} + X_i'\beta + \eta_t + \varepsilon_{ijct}$$

where $Exposure_{ct}$ is a measure of the number of fatalities in the vicinity of court c in a given time window. All other variables are defined as in equation (2). The association between exposure and bias is captured by θ_3, while α_3 now captures the "baseline bias," i.e., the level of bias when exposure to violence is zero.[15]

Column 1 of Table 3 shows the results from estimating this equation for the conflict period, using as our exposure variable the number of fatalities in the vicinity of the court in the year preceding the trial (divided by 100 for ease of exposition). The estimate in the bottom row suggests a strong and highly significant relationship. An additional fatality (in the natural area of the court) is associated with a 0.618 percentage-point increase in bias. Baseline bias is

[15] We continue to cluster standard errors by judge. Clustering by geographical areas (a natural alternative given that exposure varies by region) yields significantly smaller standard errors than clustering by judge. We adopt the more conservative approach.

estimated at 0.135 (second row). This can be compared to the overall level of bias during this period, which is estimated at 0.192 (column 3 of Table 2). A qualitatively similar pattern emerges at the sub-district and district levels (columns 4 and 7 of Table 3).

[Table 3]

These results suggest a strong short-run effect of violence on bias. How persistent is this local effect of violence? To address this question we need to compare (a) the association between exposure to violence during a given time interval and bias in a period immediately following it and (b) the association between exposure to violence during that same interval and bias in a later period. To that end, we use as our measure of exposure the cumulative number of fatalities in the vicinity of the court during the first three years of the conflict (2000–2003). This measure can potentially affect judicial decisions both in the last year of the conflict (2004) and in the post-conflict period (2007–2010).[16]

The results for exposure at the natural-area level are in columns 2 and 3. In column 2, the sample includes cases from the last year of the conflict. In column 3 the sample covers the post-conflict period. The estimates indicate that 2000–2003 fatalities are positively associated with bias in decisions taken both in 2004 and in 2007–2010. However, the point estimate is smaller for cases decided in the post-conflict period. An additional fatality in the natural area of the court during the first three years of the conflict is associated with an increase in bias of 0.467 percentage points in the final year of conflict and of 0.280 percentage

[16] Since we have court fixed effects we cannot identify θ_0 from equation (3). We also cannot identify θ_2 for the 2004 sample since Arab judges in this sample did not move between natural areas.

points in the post-conflict period.[17] A similar pattern is observed when examining exposure at the sub-district level (columns 5–6) and at the district level (columns 8–9). Note also that, regardless of the geographical level of aggregation, the estimated bias in courts that saw no fatalities in 2000–2003 (second row) appears to be slightly higher in the post-conflict period than in the last year of the conflict. In other words, the results seem to suggest an increase in baseline bias.

Table 4 examines the association between judicial decisions in the post-conflict period and court exposure during the *entire* conflict period (2000–2004).[18] As seen in the bottom row, the association between exposure and bias is always positive. For instance, the point estimate in the first column implies that an additional fatality (in the natural area of the court) during the entire conflict period is associated with a 0.256 percentage point higher bias post-conflict.

[Table 4]

While our estimates of θ_3 suggest a positive association between judicial bias and past exposure of the court to violence, one needs to be careful before interpreting the association as causal. Much of the analysis in Tables 3-4 relies on cross-sectional variation in exposure. This raises the concern that other factors

---

[17] A formal test cannot reject the equality of these two coefficients. The same is true for other differences across periods highlighted in Table 3 (both in the baseline bias and in the association with exposure).

[18] Because there was some violence during the post-conflict period, the regressions in Table 4 also control for recent fatalities in the vicinity of the court (fully interacted with the ethnicity variables). This does not change the main results. We did not include these controls in Table 3 to preserve the same functional form for the 2004 and 2007-10 samples.

that vary at the court level might be driving the results. For example, in our data more populous cities suffer from a larger number of fatalities. If larger cities are more biased, this would lead us to over-estimate θ_3. Table 5 examines this possibility for various traits of the towns in which the courts are located (e.g. population, education and income).[19] We augment the specification of Table 4 with interactions between town traits and the set of ethnicity indicators:

$$
\begin{aligned}
(4)\qquad y_{ijct} &= \alpha_0 + \alpha_1 ArabPlaintiff_i + \alpha_3 ArabPlaintiff * ArabJudge_i \\
&+ \theta_1 Exposure_c * ArabPlaintiff_i + \theta_2 Exposure_c * ArabJudge_i \\
&+ \theta_3 Exposure_c * ArabPlaintiff * ArabJudge_i \\
&+ \mu_1 TownTrait_c * ArabPlaintiff_i + \mu_2 TownTrait_c * ArabJudge_i \\
&+ \mu_3 TownTrait_c * ArabPlaintiff * ArabJudge_i \\
&+ \delta_c + \gamma_j + \lambda tenure_{jt} + X_i'\beta + \eta_t + \varepsilon_{ijct}
\end{aligned}
$$

For each of the town traits, Table 5 reports the estimated θ_3 and μ_3, where the exposure variable is the number of fatalities in the natural area around the court during the conflict. For ease of comparison, the first column replicates the baseline results of column 1 in Table 4.

[Table 5]

We begin (in column 2) by controlling for whether the town is ethnically mixed, i.e. has significant numbers of both Arab and Jewish residents.[20] Whether ethnic bias differs across ethnically mixed and ethnically homogeneous towns is

[19] Data on these traits come from the Israeli Central Bureau of Statistics. We use data from 2006 but the town traits vary little over time.

[20] More than two thirds of Arab Israelis live in strictly Arab localities. Most of the rest live in localities defined by the CBS as "mixed": localities in which the majority of the population is Jewish, but which have a significant minority of Arab residents (e.g. Jerusalem).

an interesting question in its own right (beyond testing if the results in Table 4 might be driven by this town trait). In particular, the "contact hypothesis" (Allport 1954) argues that contact between groups should, under certain conditions, reduce tensions between them. It is well recognized, however, that simply bringing the groups into closer proximity with one another is no guarantee of social harmony and might even worsen the situation by making the groups psychologically more salient (Brewer and Brown 1998).[21] Similarly, Kaufmann (1996) argues that localities with mixed ethnic populations create security dilemmas that intensify violence. Of the twenty five courts we study, five are located in mixed towns. The results in column 2 indicate that these mixed towns are indeed associated with lower bias, but the association is not statistically different from zero. At the same time, the estimate of θ_3 remains positive and statistically significant.

For the remaining town traits, we construct indicators that split the towns at the median value of each trait in the sample. We look at population size; mean per capita income; education (measured by the high school completion rate); the proportion of families with four children or more (this might be viewed as a crude proxy for religiosity); and the median age. Results are in columns 3-7. In all cases, the association between exposure to violence and bias maintains its size and statistical significance (if anything, the estimated θ_3 tends to increase). Interestingly, except for population size, none of the traits appears to be systematically associated with ethnic bias. As for population size

[21] Indeed, in his study of ethnic discrimination in the Israeli used cars market, Zussman (2013) finds that Jewish car sellers residing in mixed localities tend to discriminate more than others against Arab buyers.

(column 3), we find that controlling for exposure to violence, larger towns exhibit less bias.

The results in Table 5 notwithstanding, one might still worry that fatal terrorist attacks occur more often in regions that are characterized by greater bias to begin with. This is unlikely since, as explained earlier, the violence is not initiated by the local population and is motivated by national rather than local considerations. Even so, this concern is worth addressing directly. One way to do this would be to examine the relationship between judicial decisions made in the *pre-conflict* period and our measure of court exposure to violence during the conflict. Unfortunately, the court documents required for the analysis became available online only in late 2000 (after the conflict started). Nevertheless, we can employ a similar approach using data on judicial decisions that were made before the conflict was over. These data allow us to examine whether judicial bias at the court in the early years of the conflict is associated with exposure of the court to violence in the last year of the conflict.

Specifically, we estimate the full regression from Table 4, using fatalities from the last year of the conflict as our measure of court exposure to violence. We do this separately for two samples of judicial decisions: one from the early years of the conflict (2000–2003) and the other from the post-conflict period (2007–2010). This is done for all three geographical areas. Results are in Appendix H, Table H1 and support the causal interpretation of the association between violence and bias. Court exposure in the last year of the conflict is always positively associated with judicial bias in the post-conflict period (although the association is statistically insignificant in one case). In contrast, ethnic bias in 2000–2003 is never positively associated with the court's exposure to violence in 2004.

Another way to address the concern of spurious correlation uses variation in violence intensity throughout the conflict period. Specifically, we examine whether ethnic bias during the conflict period is associated with fatalities in the vicinity of the court in the year *following* the trial as well as in the year *preceding* the trial. The full results are reported in Appendix H, Table H2. For all geographical regions, bias is strongly and positively associated with past fatalities and not at all associated with future fatalities. Thus, there is again no evidence to suggest that bias-prone courts were exposed to disproportionately more violence.

To sum up, during the conflict period ethnic bias in judicial decisions is strongly correlated with exposure to violence. This association weakens but persists into the post-conflict period. Taken together with the results of the previous section, which show that the overall level of bias does not decline, this pattern suggests a convergence across courts toward a relatively high level of bias in the post-conflict period.

6. Channels of persistence

We have established the existence of a positive relationship between ethnic bias and past exposure to violence at the court level. But what drives this persistent effect? One possibility is that ethnic violence produces enduring changes in exposed courts due to local dynamics, such as ethnic tensions in the locality where the court is situated or in the court itself. Alternatively, violence may have a persistent effect on the legal system via the personal experience and memories of individual judges, regardless of where they currently work. Knowing the relative importance of the two channels is important for designing effective post-conflict policies. In this section we attempt to disentangle these two mechanisms by exploiting the fact that judges do not stay in the same location throughout the period under investigation.

As detailed in Appendix A, we collect data on the number of fatalities each judge was exposed to in her place of employment (either as a judge or in different occupations) in each month of the conflict. From these data we construct several measures of personal exposure. These measures are all based on the number of fatalities in the natural area surrounding the judge's place (or places) of employment during the conflict period.[22]

Again, it is instructive to start by examining the raw data. In Figure 5 we distinguish between cases handled by judges with "low personal exposure" and judges with "high personal exposure". The former category consists of judges who worked throughout the conflict period in natural areas that saw no fatalities, while the latter consists of judges who worked in areas with a positive number of fatalities. The first bar in each category shows the simple difference-

[22] In principle, it could be useful to also know the judge's exposure at her place of residence and whether relatives and friends of the judge were killed or injured by ethnic violence. Such information is difficult to collect since a judge's place of residence is confidential (surveying judges requires approval by the chief justice of the Supreme Court). However, this difficulty does not seem to be a critical concern for the present analysis. First, if many judges reside in the area where they work (which seems likely), then exposure to violence at the workplace would be highly correlated with exposure to violence at the place of residence. Second, the likelihood of a civilian in Israel being killed in a terrorist attack during the conflict period was less than 1 in 10,000. This suggests that very few (if any) of the judges in our sample saw a family member killed. Exposure to terrorism of the wider set of acquaintances is probably well-proxied by the exposure variables we use.

in-differences estimate of ethnic bias during the conflict period. While ethnic bias is 8% in cases ruled by judges with low personal exposure, it is 21% for judges with high personal exposure. Note, however, that since the judge's place of employment is for the most part the court in which she rules, it is difficult to say whether this difference is due to personal exposure or to court exposure. The correlation between this measure of personal exposure and the analogous indicator of positive court exposure is 0.77. In the post-conflict period, however, cases are often handled by judges who were not working in the same region during the conflict period. As a result, the correlation between the personal and court exposure indicators drops to 0.33. This allows us to get a sense of the separate effects of court and personal exposure.

[Figure 5]

Consider post-conflict cases handled by judges with low personal exposure, and now distinguish between them according to court exposure during the conflict (no fatalities versus a positive number of fatalities). Ethnic bias is only 2% in the cases ruled in low violence courts while it is 20% in cases ruled in high violence courts (note however the wide confidence intervals). Remarkably, the pattern is similar for cases ruled by judges with high personal exposure, with bias of 5% in low violence courts and 27% in high-violence courts. There seems to be more ethnic bias (20%) in decisions made in high-violence courts by judges with low personal exposure, than in decisions made in low-violence courts by judges with high personal exposure (5%). In other words, court exposure seems to matter much more than personal exposure: where you stand depends on where you sit.

To examine this more carefully, we estimate equation (3) for the post-conflict period, replacing the court-level measures $Exposure_{ct}$ with personal-level measures. Results are reported in Table 6. Notice that since we do not have personal exposure data for all judges, the sample size falls from 1,405 to 1,322 cases.[23] To facilitate comparison to the court-level analysis, in column 1 we replicate the regression of column 1 in Table 4 for the restricted sample. The results are similar to the original ones.

[Table 6]

Our first measure of personal exposure is the mean monthly exposure to fatalities in the natural area of the judge's place of employment during the entire conflict period. As seen in column 2, and consistent with the patterns in Figure 5, the association between this measure of personal exposure and post-conflict judicial bias is weak and statistically insignificant.

One possible explanation for this surprisingly weak association is that we do not measure correctly the way judges recall past experiences. In other contexts it has been argued that evaluations of past aversive personal experiences are dominated by the affect associated with the *worst* and the *final* moments of the episodes – not with the average or cumulative experience (this is known as the peak-end rule, see e.g., Kahneman et al. 1993; Redelmeier and Kahneman 1996). It is conceivable that a similar mechanism is operating in the present context. We therefore define *peak exposure* as the exposure at the worst month of the conflict, i.e., the maximal monthly number of fatalities in the natural area of the judge's place of employment during the conflict period. We

[23] Further excluding from the analysis judges with incomplete employment history does not affect the results.

define *late exposure* as the mean monthly number of fatalities in the natural area of the judge's place of employment during the last year of the conflict (2004). As seen in column 3, peak exposure is not associated with bias in the post-conflict period. Late exposure does seem to be positively associated with bias, but the association is not statistically different from zero (column 4).

Finally, we attempt to disentangle the two mechanisms, one operating at the local (court) level and the other operating at the individual (judge) level. In columns 5–7 we include both the court exposure measure and each of the personal exposure measures.[24] Consistent with the patterns observed in Figure 5, we find that controlling for court exposure, none of the personal exposure measures is associated with judicial bias in the post-conflict period. At the same time, the court exposure measure itself remains positively associated with judicial bias.

A possible difficulty in interpreting the results in Table 6 is that judge movements may not be random. For example, it is conceivable that over time judges move to the vicinity of their home town or, more generally, closer to the place they care about. Thus, court exposure in fact captures violence at the place the judge cared about even during the conflict. To address this issue, in Appendix I we re-examine the association between judicial bias and violence during the conflict period. We construct a new measure of personal exposure to violence. For each judge we locate the natural area in which she spent the most time during the *post-conflict* period. We next compute for each decision the number of fatalities in the preceding year in that natural area (i.e. in the area where the

[24] The simple correlation coefficient between court exposure and the three measures of personal exposure ranges from 0.53 to 0.69.

judge will work after the conflict). This measure is available for 1,583 of the 1,748 cases in the conflict period. Finally, we augment the baseline specification for the conflict period (column 1 of Table 3) with this new measure and its interactions with the ethnicity indicators. In other words, we run a "horse race" between the original court exposure variable and the new personal exposure measure. If judges care about the place they work in during the post-conflict period, the new measure should be positively associated with ethnic bias, while court exposure should not. Instead, we find that the estimated effect of court exposure maintains its size and statistical significance, while the new personal exposure variable is not associated with ethnic bias. This seems to rule out the alternative interpretation of Table 6.

7. Conclusion

This paper sought to contribute to our understanding of the effects of ethnic conflict on intergroup relations and on the functioning of legal institutions. We employed a decision-based methodology to measure the extent of ethnic bias among an elite population that arguably should be most immune to ethnic biases: judges.

Looking at the country as a whole, we uncover a striking fact: despite a dramatic drop in the number of fatalities in the years that followed the Second Intifada, the level of ethnic bias remained roughly similar to that observed during the height of the conflict.

Our second main finding leverages exogenous variation in the intensity of ethnic violence. The analysis uncovers a persistent adverse effect of exposure to violence on judicial decisions. In the post-Intifada years, courts that had experienced more violence still tended to exhibit greater ethnic bias than courts situated in quieter regions.

Perhaps most importantly, we find that the persistent effect of violence operates primarily at the local (court) level rather than at the personal (judge) level. Individuals who were personally exposed to high levels of violence during the conflict show little systematic bias in subsequent years if located in a relatively peaceful environment. At the same time, individuals located in afflicted areas display high levels of bias even if they were not personally exposed to high levels of violence during the conflict.

References

Abrams, David S., Marianne Bertrand, and Sendhil Mullainathan. 2012. "Do Judges Vary in Their Treatment of Race?" *Journal of Legal Studies*, 41(2): 347–383.

Acharya, Avidit, Matthew Blackwell, and Maya Sen. 2015. "The Political Legacy of American Slavery." *Unpublished manuscript*.

Alesina, Alberto, Paola Giuliano, and Nathan Nunn. 2013. "On the Origins of Gender Roles: Women and the Plough." *Quarterly Journal of Economics*, 128 (2): 469–530.

Alesina, Alberto, and Nicola Fuchs-Schündeln. 2007. **"**Goodbye Lenin (or Not?): The Effect of Communism on People." *American Economic Review*, 97(4): 1507–1528.

Alesina, Alberto, and Eliana La Ferrara. 2014. "A Test of Racial Bias in Capital Sentencing." *American Economic Review*, 104(11): 3397-3433.

Allport, Gordon W. 1954. *The Nature of Prejudice*. Reading, MA: Addison-Wesley.

Anwar, Shamena, Patrick Bayer, and Randi Hjalmarsson. 2012. "The Impact of Jury Race in Criminal Trials." *Quarterly Journal of Economics*, 127(2): 1017–1055.

Bellows, John, and Edward Miguel. 2009. ''War and Collective Action in Sierra Leone.'' *Journal of Public Economics*, 93 (11–12): 1144–57.

Berrebi, Claude, and Esteban F. Klor. 2008. "Are Voters Sensitive to Terrorism? Direct Evidence from the Israeli Electorate." *American Political Science Review*, 102(3): 279-301.

Bertrand, Marianne, and Esther Duflo. Forthcoming. "Field Experiments on Discrimination." in Abhijit Banerjee and Esther Duflo eds., *Handbook of Field Experiments.*

Besley, Timothy, and Marta Reynal-Querol. 2014. "The Legacy of Historical Conflict: Evidence from Africa." *American Political Science Review*, 108(2): 319-336.

Besley, Timothy, and Torsten Persson. 2011. *Pillars of Prosperity: The Political Economics of Development Clusters*. Princeton, NJ: Princeton University Press.

Blattman, Christopher. 2009. "From Violence to Voting: War and Political Participation in Uganda." *American Political Science Review*, 103(2): 231-247.

Bornstein, Gary. 2003. "Intergroup Conflict: Individual, Group, and Collective Interests." *Personality and Social Psychology Review*, 7(2): 129–145.

Brewer, Marilynn, and Rupert Brown. 1998. "Intergroup relations." In: Gilbert, Daniel, Susan Fiske and Gardner Lindzey (editors), *The Handbook Of Social Psychology*, Vol. II, New York: McGraw-Hill, pp. 554-594.

Callen, Michael, Mohammad Isaqzadeh, James D. Long, and Charles Sprenger. 2014. "Violence and Risk Preference: Experimental Evidence from Afghanistan." *American Economic Review*, 104(1): 123-48.

Canetti-Nisim, Daphna, Eran Halperin, Keren Sharvit, and Stevan E. Hobfoll. 2009. "A New Stress-Based Model of Political Extremism: Personal Exposure to Terrorism, Psychological Distress, and Exclusionist Political Attitudes." *Journal of Conflict Resolution*, 53(3): 363–89.

Cassar, Alessandra, Pauline Grosjean and Sam Whitt. 2013. "Legacies of Violence: Trust and Market Development." *Journal of Economic Growth*, 18(3): 285-318.

Chandra, Kanchan (editor). 2012. *Constructivist Theories of Ethnic Politics*. New York: Oxford University Press.

Chen, Yan, and Sherry Xin Li. 2009. "Group Identity and Social Preferences." *American Economic Review*, 99(1): 431–457.

Cilliers, Jacobus, Oeindrila Dube and Bilal Siddiqi. 2015. "Can the Wounds of War be Healed? Experimental Evidence from Reconciliation in Sierra Leone". *Unpublished manuscript*.

Collier, Paul, Lani Elliot, Havard Hegre, Anke Hoeffler, Marta Reynal-Querol, and Nicholas Sambanis. 2003. *Breaking the Conflict Trap: Civil War and Development Policy*. Washington DC: The World Bank.

Fearon, James, Macartan Humphreys, and Jeremy Weinstein. Forthcoming. "How Does Development Assistance Affect Collective Action Capacity? Results from a Field Experiment in Post-Conflict Liberia." *American Political Science Review*.

Fong, Christina M., and Erzo F. P. Luttmer. 2009. "What Determines Giving to Hurricane Katrina Victims? Experimental Evidence on Racial Group Loyalty." *American Economic Journal: Applied Economics*, 1(2): 64–87.

Getmansky, Anna and Thomas Zeitzoff. 2014. "Terrorism and Voting: The Effect of Rocket Threat on Voting in Israeli Elections." *American Political Science Review*, 108(3): 588-604.

Gilligan, Michael, Benjamin Pasquale, and Cyrus Samii. 2014. "Civil War and Social Cohesion: Lab-in-the-Field Evidence from Nepal." *American Journal of Political Science*, 58(3): 604–619.

Glaeser, Edward L., and Bruce Sacerdote. 2003. "Sentencing in Homicide Cases and the Roles of Vengeance." *Journal of Legal Studies*, 32(2): 363–382.

Gould, Eric D., and Esteban F. Klor. 2010. "Does Terrorism Work?" *Quarterly Journal of Economics*, 125(4): 1459–1510.

Grossman, Guy, Oren Gazal-Ayal, Samuel Pimentel and Jeremy Weinstein. 2016. "Descriptive Representation and Judicial Outcomes in Multi-Ethnic Societies." *American Journal of Political Science*, forthcoming.

Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2007. "Why Does Ethnic Diversity Undermine Public Goods Provision?" *American Political Science Review*, 101(4): 709-725.

Hirsch-Hoefler, Sivan, Daphna Canetti, Carmit Rapaport and Stevan Hobfoll. Forthcoming. "Conflict will Harden your Heart: Exposure to Violence, Psychological Distress, and Peace Barriers in Israel and Palestine". *British Journal of Political Science*.

Kahneman, Daniel, Barbara L. Fredrickson, Charles A. Schreiber, and Donald A. Redelmeier. 1993. "When More Pain is Preferred to Less: Adding a Better End." *Psychological Science*, 4(6): 401–405.

Kaufmann, Chaim. 1996. "Possible and Impossible Solutions to Ethnic Civil Wars." *International Security*, 20(1): 136–175.

Klor, Esteban F., and Moses Shayo. 2010. "Social Identity and Preferences Over Redistribution." *Journal of Public Economics*, 94(3-4): 269-278.

Laitin, David. 2007. *Nations, States, and Violence*. Oxford: Oxford University Press.

Lyall, Jason, Graeme Blair, and Kosuke Imai. 2013. "Explaining Support for Combatants During Wartime: A Survey Experiment in Afghanistan." *American Political Science Review*, 107(4): 679–705.

Neumark, David. Forthcoming. "Experimental Research on Labor Market Discrimination." *Journal of Economic Literature.*

Nunn, Nathan, and Leonard Wantchekon. 2011. “The Slave Trade and the Origins of Mistrust in Africa.” *American Economic Review*, 101(7): 3221–3252.

Redelmeier, Donald A., and Daniel Kahneman. 1996. “Patient’s Memories of Painful Medical Treatments: Real-time and Retrospective Evaluations of Two Minimally Invasive Procedures.” *Pain*, 66(1): 3-8.

Rohner, Dominic, Mathias Thoenig and Fabrizio Zilibotti. 2013. “Seeds of Distrust: Conflict in Uganda.” *Journal of Economic Growth,* 18(3): 217-252.

Sambanis, Nicholas, and Moses Shayo. 2013. **“**Social Identification and Ethnic Conflict.” *American Political Science Review,* 107(2): 294–325.

Shayo, Moses. 2009. “A Model of Social Identity with an Application to Political Economy: Nation, Class and Redistribution.” *American Political Science Review,* 103 (2): 147–174.

Shayo, Moses, and Asaf Zussman. 2011. “Judicial Ingroup Bias in the Shadow of Terrorism.” *Quarterly Journal of Economics*, 126(3): 1447–1484.

Tajfel, Henri, Michael G. Billig, Robert P. Bundy, and Claude Flament. 1971. “Social Categorization and Intergroup Behavior.” *European Journal of Social Psychology*, 1(2): 149–178.

Voigtländer, Nico, and Hans-Joachim Voth. 2012. “Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany.” *Quarterly Journal of Economics*, 127(3): 1339–1392.

Voors, Maarten J., Eleonora E. M. Nillesen, Philip Verwimp, Erwin H. Bulte, Robert Lensink, and Daan P. Van Soest. 2012. “Violent Conflict and Behavior: A Field Experiment in Burundi.” *American Economic Review*, 102(2): 941–964.

Zussman, Asaf. 2013. “Ethnic Discrimination: Lessons from the Israeli Online Market for Used Cars.” *Economic Journal*, 123(572): F433–F468.

TABLE 1: CASES BY ETHNICITY OF JUDGE AND LITIGANTS

			Period		
Judge	Plaintiff	Defendant	Conflict	Post-conflict	Total
Arab	Arab	Jewish	280	175	455
			(16.02%)	(12.46%)	(14.43%)
	Jewish	Arab	265	180	445
			(15.16%)	(12.81%)	(14.11%)
Jewish	Arab	Jewish	485	396	881
			(27.75%)	(28.19%)	(27.94%)
	Jewish	Arab	718	654	1,372
			(41.08%)	(46.55%)	(43.51%)
Total			1,748	1,405	3,153
			(100%)	(100%)	(100%)

TABLE 2: ETHNIC BIAS DURING AND AFTER THE CONFLICT

Dependent variable: claim accepted

	Conflict	Post-conflict	Conflict	Post-conflict
	(1)	(2)	(3)	(4)
Arab plaintiff	-0.151***	-0.169***	-0.117***	-0.170***
	(0.026)	(0.033)	(0.031)	(0.044)
Arab judge	-0.077*	-0.005		
	(0.044)	(0.031)		
Arab judge*Arab plaintiff	0.170***	0.177***	0.192***	0.179***
	(0.054)	(0.046)	(0.049)	(0.057)
Court FEs	Yes	Yes	Yes	Yes
Judge FEs & tenure	No	No	Yes	Yes
Case characteristics	No	No	Yes	Yes
Time controls	No	No	Yes	Yes
Observations	1,748	1,405	1,748	1,405
R-squared	0.044	0.057	0.248	0.258

Notes: Regressions are estimated by OLS. Standard errors, clustered by judge, are in parentheses. Case characteristics include: number of plaintiffs; number of defendants; share of private plaintiffs; share of private defendants; share of male plaintiffs; share of male defendants; monetary compensation requested (and an indicator for missing values); indicators for claim subjects; an indicator for "defense present"; and an indicator for cases where the defendant filed a counterclaim. Time controls include indicators for year, month, and day of week.

*, **, *** represent statistical significance at the 10, 5, and 1 percent levels.

TABLE 3: EFFECTS OF VIOLENCE IN THE SHORT AND THE LONG RUN

Dependent variable: claim accepted

	Natural Area			Sub-district			District		
Cases from:	2000–2004	2004	2007–2010	2000–2004	2004	2007–2010	2000–2004	2004	2007–2010
Court exposure during:	Preceding year	2000–2003	2000–2003	Preceding year	2000–2003	2000–2003	Preceding year	2000–2003	2000–2003
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Arab plaintiff	-0.096***	-0.077	-0.162***	-0.095***	-0.059	-0.151**	-0.084*	-0.031	-0.141
	(0.032)	(0.080)	(0.056)	(0.034)	(0.084)	(0.062)	(0.045)	(0.135)	(0.092)
Arab judge*Arab plaintiff	0.135***	0.102	0.131*	0.121**	0.017	0.106	0.112*	-0.020	0.061
	(0.044)	(0.128)	(0.071)	(0.049)	(0.152)	(0.086)	(0.060)	(0.194)	(0.117)
Court exposure	0.044			0.041			0.067		
	(0.152)			(0.150)			(0.111)		
Arab plaintiff*Court exposure	-0.143	-0.219**	-0.021	-0.136	-0.238**	-0.041	-0.133	-0.176	-0.036
	(0.134)	(0.100)	(0.080)	(0.135)	(0.102)	(0.086)	(0.138)	(0.135)	(0.090)
Arab judge*Court exposure	-0.190		2.326	-0.249		2.088*	-0.300		
	(0.335)		(2.436)	(0.313)		(1.131)	(0.191)		
Arab plaintiff*Arab judge*Court exposure	0.618***	0.467**	0.280**	0.627***	0.601***	0.249	0.436**	0.410**	0.219
	(0.178)	(0.193)	(0.124)	(0.185)	(0.227)	(0.154)	(0.171)	(0.196)	(0.151)
Observations	1,748	589	1,405	1,748	589	1,405	1,748	589	1,405
R-squared	0.250	0.328	0.260	0.249	0.329	0.260	0.249	0.326	0.259

Notes: Analysis includes cases from the period indicated in the first row of the column title. Court exposure is the cumulative number of civilian fatalities in the vicinity (natural area/sub-district/district) of the court during the period indicated in the second row of the column title. Fatality figures are divided by 100 for clarity. Regressions are estimated by OLS. Standard errors, clustered by judge, are in parentheses. All regressions include court fixed effects and the full set of judge, case and time controls from columns 3-4 of Table 2. *Arab judge*Court exposure* is omitted from column 9 because two districts had the same absolute number of fatalities during the conflict, which generates perfect multicollinearity.

*, **, *** represent statistical significance at the 10, 5, and 1 percent levels.

TABLE 4: LONG-RUN EFFECTS OF EXPOSURE TO VIOLENCE – COURT LEVEL ANALYSIS

Cases from the post-conflict period (2007–2010)

Dependent variable: claim accepted

	Natural Area	Sub-district	District
	(1)	(2)	(3)
Arab plaintiff	-0.157***	-0.150**	-0.148*
	(0.056)	(0.062)	(0.088)
Arab judge*Arab plaintiff	0.146**	0.102	0.034
	(0.066)	(0.085)	(0.108)
Arab plaintiff*Court exposure	-0.032	-0.040	-0.033
	(0.074)	(0.083)	(0.083)
Arab judge*Court exposure	2.322	2.127*	
	(2.412)	(1.167)	
Arab plaintiff*Arab judge*Court exposure	0.256**	0.254*	0.232*
	(0.108)	(0.150)	(0.139)
Observations	1,405	1,405	1,405
R-squared	0.267	0.264	0.265

Notes: Court exposure is the cumulative number of civilian fatalities in the vicinity (natural area/sub-district/district) of the court during the conflict period (28/9/2000–31/12/2004), divided by 100 for clarity. Regressions are estimated by OLS. Standard errors, clustered by judge, are in parentheses. All regressions include court fixed effects and the full set of judge, case and time controls from columns 3-4 of Table 2. Regressions also control for the number of civilian fatalities (separately from the Israeli-Palestinian conflict and the 2006 Lebanon War) in the vicinity of the court in the year preceding the trial fully interacted with indicators for judge and plaintiff ethnicity. *Arab judge*Court exposure* is omitted from column 3 because two districts had the same absolute number of fatalities during the conflict, which generates perfect multicollinearity.

*, **, *** represent statistical significance at the 10, 5, and 1 percent levels.

TABLE 5: LONG-RUN EFFECTS OF EXPOSURE TO VIOLENCE, CONTROLLING FOR TOWN TRAITS

Cases from the post-conflict period (2007–2010)

Dependent variable: claim accepted

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Arab plaintiff	-0.157***	-0.158***	-0.172***	-0.187**	-0.193**	-0.155***	-0.159**
	(0.056)	(0.045)	(0.066)	(0.083)	(0.081)	(0.055)	(0.071)
Arab judge*Arab plaintiff	0.146**	0.147**	0.259***	0.156	0.164*	0.049	0.154**
	(0.066)	(0.061)	(0.066)	(0.095)	(0.093)	(0.151)	(0.077)
Arab plaintiff*Arab judge*Court exposure	0.256**	0.423***	0.365***	0.240***	0.245***	0.330**	0.318***
	(0.108)	(0.144)	(0.092)	(0.088)	(0.087)	(0.127)	(0.114)
Arab plaintiff*Arab judge*Mixed ethnicity town		-0.222					
		(0.165)					
Arab plaintiff*Arab judge*Large town			-0.295***				
			(0.103)				
Arab plaintiff*Arab judge*High-income town				0.007			
				(0.113)			
Arab plaintiff*Arab judge*High-education town					-0.006		
					(0.111)		
Arab plaintiff*Arab judge*Large-family town						0.103	
						(0.154)	
Arab plaintiff*Arab judge*High-median-age town							-0.089
							(0.144)
Observations	1,405	1,405	1,405	1,405	1,405	1,405	1,405
R-squared	0.267	0.268	0.272	0.268	0.268	0.268	0.267

Notes: Court exposure is the cumulative number of civilian fatalities in the natural area of the court during the conflict period divided by 100. Town traits are for 2006 and relate to the town where the court is located (source: Israeli Central Bureau of Statistics). Mixed ethnicity town is an indicator for towns with a significant numbers of both Arab and Jewish residents. Large town is an indicator for towns with a population above the median across cases (median=64,200). High-income town is an indicator for towns with an above median per-capita income (NIS 2,620). High-education town is an indicator for towns with an above median high-school matriculation rate (51.3%). Large-family town is an indicator for towns with an above median share of families with at least four children (9.35%). High-median-age town is an indicator for towns where the median age is above the median across cases (32). Regressions are estimated by OLS. Standard errors, clustered by judge, are in parentheses. All regressions include the full set of controls from column 1 of Table 4 as well as interactions between the relevant town trait and the *Arab plaintiff* and *Arab judge* indicators.

*, **, *** represent statistical significance at the 10, 5, and 1 percent levels.

TABLE 6: CHANNELS OF PERSISTENCE – PERSONAL- AND COURT-LEVEL EFFECTS OF VIOLENCE

Cases from the post-conflict period (2007–2010)

Dependent variable: claim accepted

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Arab plaintiff	-0.158***	-0.160***	-0.169***	-0.174***	-0.156***	-0.161***	-0.153**
	(0.057)	(0.051)	(0.059)	(0.046)	(0.055)	(0.061)	(0.065)
Arab judge*Arab plaintiff	0.105	0.160**	0.174**	0.169**	0.106	0.116	0.103
	(0.078)	(0.070)	(0.073)	(0.067)	(0.077)	(0.078)	(0.083)
Arab plaintiff*Arab judge*Court exposure	0.288**				0.294*	0.323**	0.307*
	(0.116)				(0.160)	(0.137)	(0.183)
Arab plaintiff*Arab judge*Mean personal exposure		0.054			-0.042		
		(0.115)			(0.108)		
Arab plaintiff*Arab judge*Peak personal exposure			0.000			-0.006	
			(0.007)			(0.007)	
Arab plaintiff*Arab judge*Late personal exposure				0.188			-0.043
				(0.118)			(0.175)
Observations	1,322	1,322	1,322	1,293	1,322	1,322	1,293
R-squared	0.268	0.266	0.266	0.266	0.268	0.268	0.268

Notes: Court exposure is the cumulative number of civilian fatalities in the natural area of the court during the conflict period (28/9/2000–31/12/2004). Fatality figures are divided by 100 for clarity. Mean personal exposure is the mean monthly number of civilian fatalities in the natural area of the judge's place of employment during the conflict period. Peak exposure is the maximal monthly number of civilian fatalities in the natural area of the judge's place of employment during the conflict period. Late exposure is the mean monthly number of civilian fatalities in the natural area of the judge's place of employment during the last year of the conflict (2004). Regressions are estimated by OLS. Standard errors, clustered by judge, are in parentheses. All regressions include the full set of controls from column 1 of Table 4.

*, **, *** represent statistical significance at the 10, 5, and 1 percent levels.

Figure 1: Civilian Fatalities by Year and District

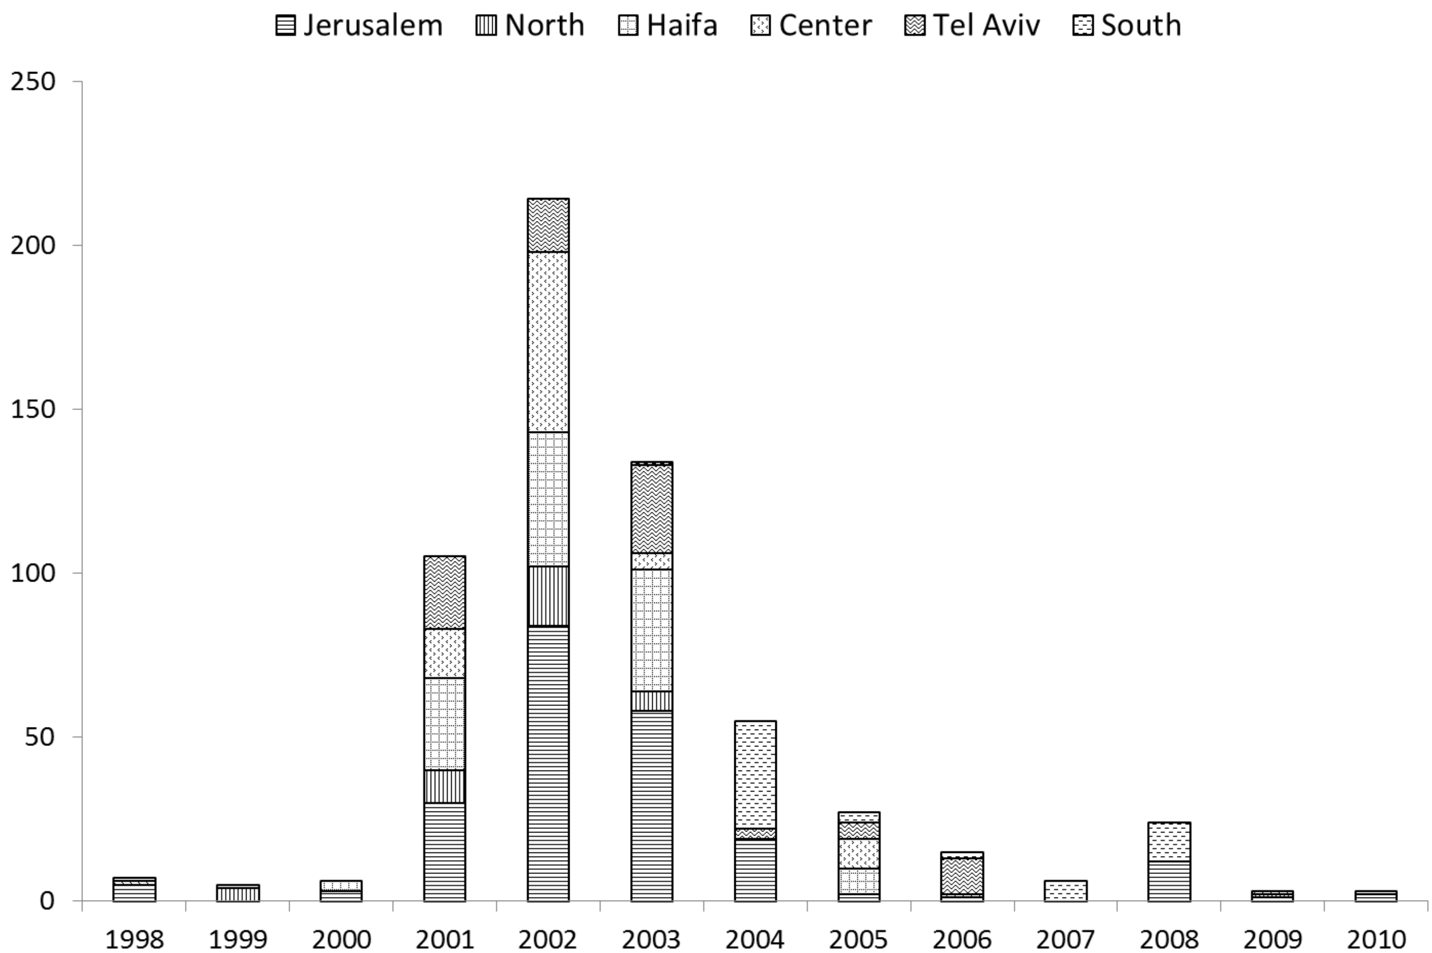

Figure 2: Ethnic Bias by Period

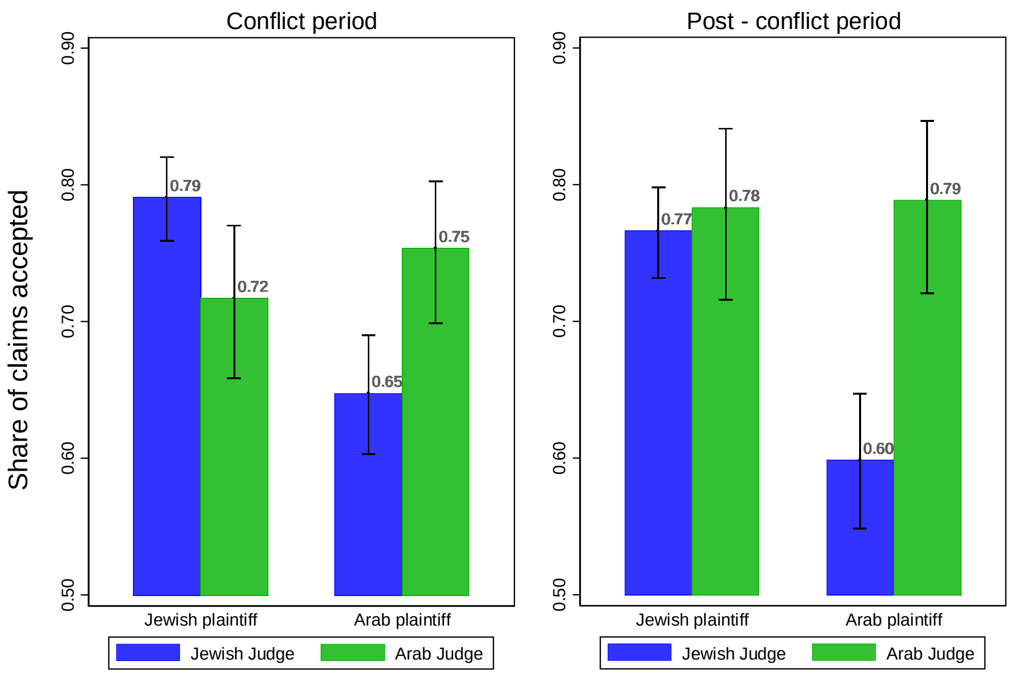

Capped ranges indicate 95% confidence intervals.

Figure 3: The Evolution of Ethnic Bias

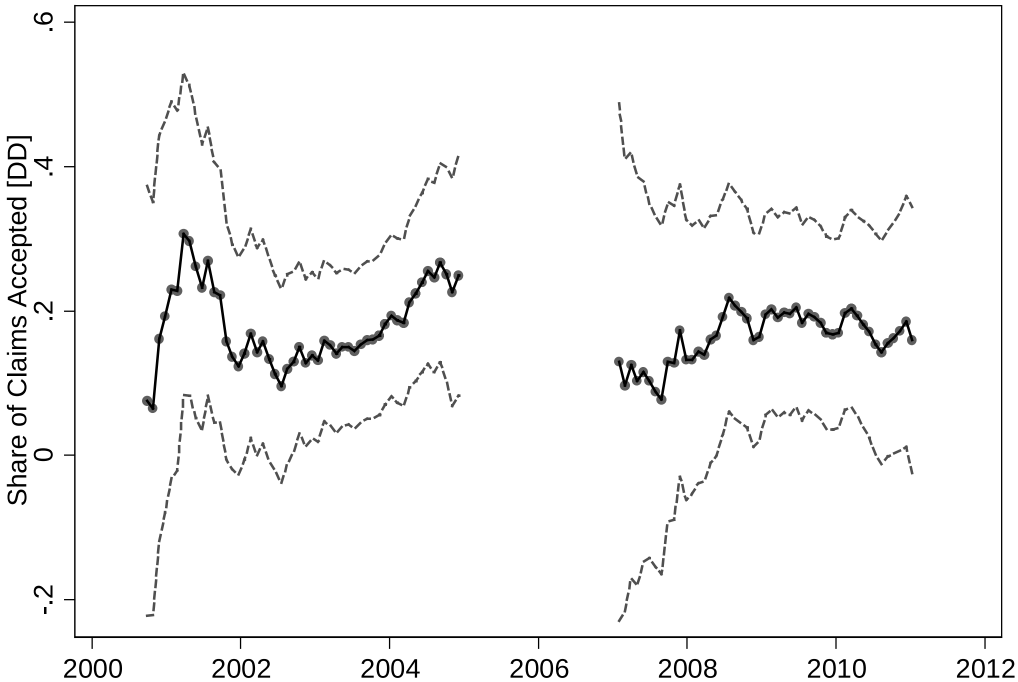

Diff-in-diff estimates in two-year moving windows with at least 100 obs (95% confidence intervals).

Figure 4: Ethnic Bias by Exposure to Violence During the Conflict

Diff-in-diff estimates of bias in two-year moving windows with at least 100 obs.
Based on 3,153 cases, 915 of which in courts that saw no fatalities during the conflict period.

Figure 5: Ethnic Bias by Personal and Court Exposure

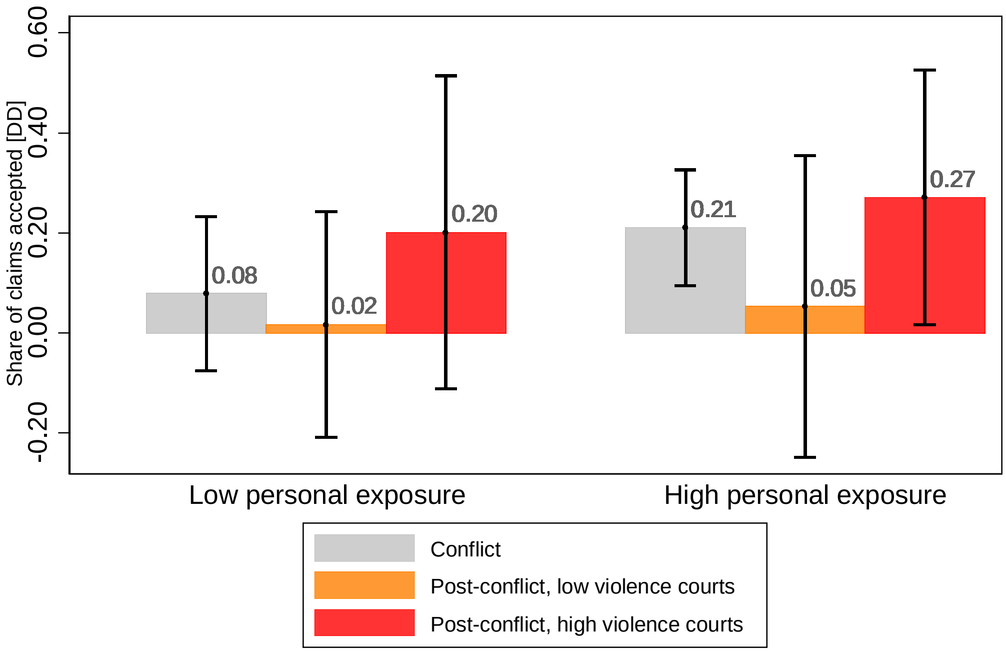

Diff-in-diff estimates. Capped ranges indicate 95% confidence intervals.
Based on 233 judges and 3,070 cases, 870 of which decided by 40 judges that saw no fatalities in their place of employment during the conflict period.