

The Impact of the Coronavirus Lockdown on Domestic Violence

Justin McCrary, *Columbia Law School*, and Sarath Sanga, *Northwestern University*

Send correspondence to: Sarath Sanga, Pritzker School of Law, Northwestern University, 375 E Chicago Ave, Chicago, IL 60611, USA; Tel: (312) 503-3100; Email: sanga@northwestern.edu

We use 911 call records and mobile device location data to study the impact of the coronavirus lockdown on domestic violence. The percent of people at home sharply increased at all hours, and nearly doubled during regular working hours, from 45% to 85%. Domestic violence increased 12% on average and 20% during working hours. Using neighborhood-level identifiers, we show that the rate of first-time abuse likely increased even more: 16% on average and 23% during working hours. Our results contribute to an urgent need to quantify the physical and psychological burdens of prolonged lockdown policies. (*JEL*: J12, D63, I18, E65)

1. Introduction

In this article, we examine the impact of the coronavirus lockdown on domestic violence. In response to the coronavirus pandemic, nearly all U.S. states issued stay-at-home orders designed to restrict the movement of people.¹ The anticipated public-health benefit of these policies was to arrest the spread of Covid-19 and lower the peak resource use of health care and emergency services. But these policies came at considerable cost. The economic

We thank Bernard Black, Vandy Howell, Katherine Litvak, Max Schanzenbach, David Schwartz, Eric Talley, Kimberly Yuracko, and seminar participants at Northwestern University for helpful comments.

1. To the best of our knowledge, Arkansas, Iowa, Nebraska, North Dakota, and South Dakota did not issue state-level orders. Our sample of 911 call records does not cover any of these states.

American Law and Economics Review
doi:10.1093/aler/ahab003

Advance Access publication on May 30, 2021

© The Author 2021. Published by Oxford University Press on behalf of the American Law and Economics Association. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

costs of lost GDP and mass unemployment are large, well-measured, and well-understood. Much less understood, however, are the physical, psychological, and emotional costs of lockdown—of which domestic violence is one tragically common example.² These costs are difficult to quantify and unlikely to be reflected in standard economic measures of welfare (Stevenson and Wolfers, 2009). To evaluate the true social cost of ongoing lockdown policies, it is therefore crucial to account for these sources of harm.³

We quantify the impact of lockdown on domestic violence by assembling a database of approximately 50 million 911 call records from 14 large U.S. cities. We also obtain mobile device location data from these same cities. Together, these data allow us to study the impacts of lockdown on at-home patterns and domestic violence. We find that both increased sharply during lockdown (Figures 1 and 2). Domestic violence calls increased throughout mid-March and peaked in early April. By the end of April, domestic violence calls returned to pre-lockdown levels. At the same time, the number of people at home also abruptly increased but stayed high throughout April. The largest increases in both at-home patterns and domestic violence occurred during weekday daytime hours, when most adults would have otherwise been at work and most children would have been in school.

While the surge in domestic violence appears to have been temporary, there are at least two reasons to suspect that the harms will be long-lasting. The first reason comes from the long-lasting nature of physical and psychological trauma. While economies can recover from large, even catastrophic, shocks, people may not be so resilient. Children who suffer or witness abuse also suffer a range of psychological, behavioral, and academic problems throughout their lives.⁴ The harms from domestic violence will therefore persist well after lockdowns end. The second reason is that violence begets violence. Inasmuch as domestic violence is state dependent, in the sense used by Heckman (1981), exposure to violence today increases

2. On the psychological costs of domestic violence, see generally, Russell (1990), White and Koss (1991), Balsam and Szymanski (2005), Truman and Morgan (2014), and Golding (1999).

3. This is especially urgent as subsequent waves of the virus—and potentially additional lockdown orders—are expected in the medium term.

4. See Edleson (1999), Kilpatrick and Williams (1997), Truman and Morgan (2014), and Snyder et al. (2016).

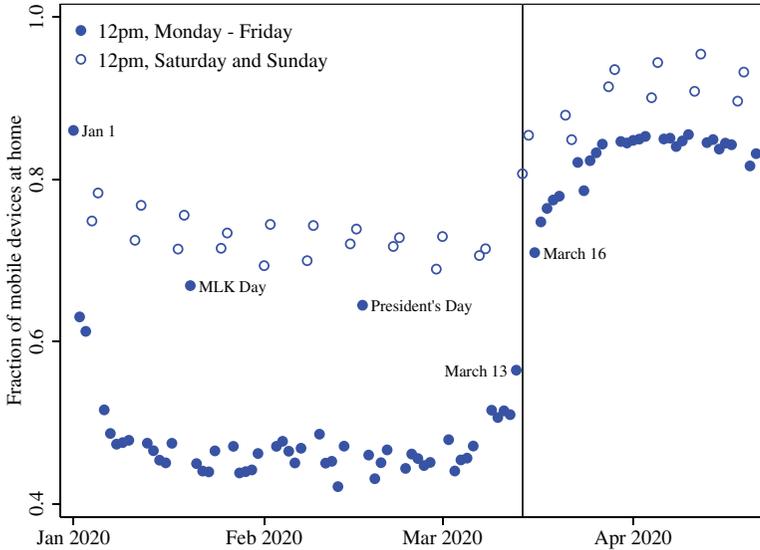


Figure 1. The impact of lockdown on being at home. The fraction of mobile devices at home at 12 p.m. each day (a proxy for the fraction of people at home) for the sample of 14 U.S. cities listed in Table 1.

the likelihood of violence tomorrow, and temporary surges in abuse will lead to a permanent increase in long-run levels.⁵

While we cannot directly estimate state dependence or the long-run impact of lockdown, we nevertheless begin to address these issues by asking whether lockdown caused people to commit abuse for the first time. The impacts of lockdown could be especially persistent if the surge in domestic violence came from households without a history of violence, or put another way, if it caused an increase on the extensive margin of domestic violence. This is because any state dependence would then act as a multiplier, causing those households to experience more violence in the future.

To estimate changes on the extensive margin, we construct neighborhood identifiers using address records attached to each 911 call. (City authorities anonymize the address to the level of about half of one city block.) We then show how the neighborhood-level data can provide a close approximation

5. Repeat-victimization in domestic violence is a fundamental concern for law enforcement authorities. See, e.g., Hanmer et al. (1999).

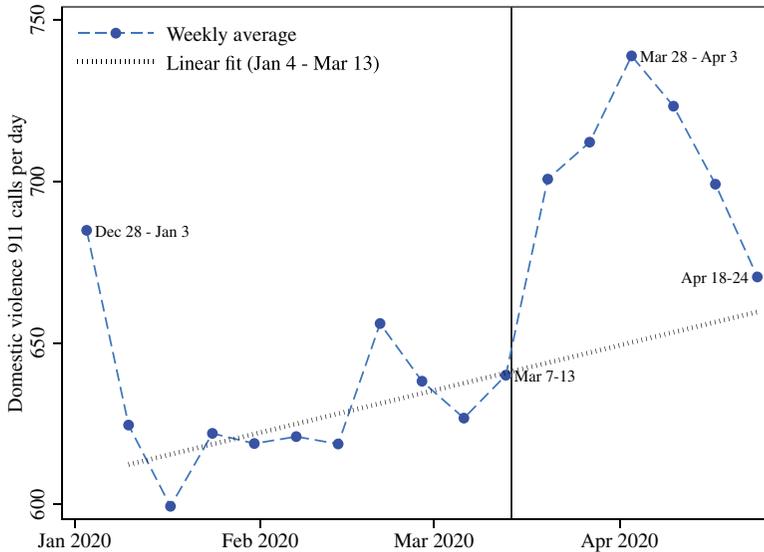


Figure 2. The impact of lockdown on domestic violence. The average daily number of domestic violence 911 calls in the sample of 14 U.S. cities listed in Table 1.

of the percentage change in the household-level extensive margin, as well as an upper bound on the household-level intensive margin. Applying our procedures to the data, we find that the impact on the extensive margin was at least twice as large as the impact on the intensive margin: 16% on the extensive margin versus an upper bound of 8% on the intensive margin. Lockdown therefore had larger impact, in percentage terms, on households without a history of domestic violence. To the extent that domestic violence is state dependent, we conclude that the temporary surge will lead to a permanent increase.

The surge in domestic violence in the wake of the coronavirus pandemic and lockdown has been documented across scores of countries. See, for example, [Perez-Vincent et al. \(2020\)](#) (Argentina); [Sifat \(2020\)](#) (Bangladesh); [Das et al. \(2020\)](#) and [Ravindran and Shah \(2020\)](#) (India); [Dahal et al. \(2020\)](#) (Nepal); [Calderon-Anyosa and Kaufman \(2020\)](#) (Peru); [Sediri et al. \(2020\)](#) (Tunisia); [Leslie and Wilson \(2020\)](#) and [Boserup et al. \(2020\)](#) (United States); [Javed and Mehmood \(2020\)](#) (various countries)—among many others. The closest study to ours is perhaps

Leslie and Wilson (2020), which also studies the U.S. context. Virtually all studies demonstrate that domestic violence significantly increased in the post-lockdown period.

Our study provides three unique contributions to this literature: (1) we assemble the largest sample of 911 calls for service (14 U.S. cities that constitute $\sim 5\%$ of the U.S. population); (2) we provide the first credible estimate of the extent to which increased time spent at home is responsible for the surge in domestic violence (specifically, an upper bound on this estimate); and (3) we theoretically motivate and demonstrate how neighborhood-level data can be used to separately identify bounds on changes in the household-level extensive and intensive margins of domestic violence.

Two caveats are in order. First, we use the term “lockdown” as a catchall for the total disruption caused by the pandemic and subsequent suspension of regular economic and social interaction. We do so with the understanding that the pandemic and the myriad public and private responses to it likely affected domestic violence through a variety of channels. In our view, it is likely that the surge in domestic violence did not come from sheltering in place per se—that is, only from spending more time at home—but rather from spending more time at home under the uniquely stressful conditions created by the pandemic. The most straightforward reason comes from Figures 1 and 2, which show that simply being at home under lockdown was not sufficient to *sustain* an increase in domestic violence: domestic violence call volumes rose with the number of people at home, but then dropped significantly throughout April even as the number of people at home remained constant.

The second caveat is that our conclusions on changes in domestic violence are based on reported incidents that result in a 911 call. Since most incidents of domestic abuse do not result in a 911 call, the data are subject to considerable reporting bias. Moreover, reporting bias may have increased or decreased during lockdown. If domestic violence incidents were less likely to be reported during lockdown, then our results would underestimate the change in domestic violence.⁶ Indeed, lockdown made it much more difficult to leave the house and go to a shelter or stay with a friend. A

6. Perhaps the strongest evidence of increased reporting bias comes from widespread reports of an increased reluctance to contact emergency medical services for both covid-19 and non-covid-19 reasons, as well as the related increase in at-home deaths.

victim unable to leave the house might reasonably conclude that involving the police would make matters worse. Even more troubling, the return to pre-lockdown domestic violence levels could reflect additional increases in reporting bias as lockdown wore on.⁷ Nevertheless, the reporting bias inherent in our study is, relative to a typical empirical study of crime, straightforward to characterize. A typical study might use crimes reported by law enforcement authorities, and therefore be subject to the reporting bias of victims, witnesses, and police. In contrast, we use 911 calls, which are typically placed by victims or witnesses. Any reporting bias would thus come directly from the victims and witnesses themselves.

The article is organized as follows. Section 2 summarizes the 911 call records and mobile device location data. Section 3 reports the main results. Section 4 reports estimates of the response to lockdown on the extensive and intensive margins of domestic violence. Section 5 concludes.

2. Data and Summary Statistics

Our analysis combines two sources of data. The first source is 911 call records for police service from 14 large U.S. cities. The second source is mobile device location histories from those same cities. This section briefly summarizes both sources of data. Additional details on data construction can be found in the [Supplementary Appendix](#).

2.1. Calls for Police Service

We obtained records of about 52 million 911 calls for police service, of which about 6.3 million were placed during the main sample period of January 1, 2020 through April 24, 2020 (Table 1). The records come from 14 cities (or more precisely, 13 cities and 1 county) with a collective population of about 17 million people. Depending on the city, there are anywhere between 2 and 15 years of call records. The records are complete for each city beginning January 1 of the starting year listed in Table 1 and running through April 24, 2020.

At-home deaths spiked 4-fold in Detroit and 6-fold in New York City (Gillum et al., 2020).

7. See, e.g., Ott (2020).

Table 1. 911 Calls for Police Service

City/County	Population	Years	911 calls placed between January 1, 2020–April 24, 2020			Neighborhoods
			Total	Domestic violence	DV share	
Los Angeles, CA	3,979,576	2010–2020	937,629	37,009	0.04	
Chicago, IL	2,693,976	2017–2020	1,029,862	11,512	0.01	21,429
Phoenix, AZ	1,680,992	2016–2020	606,439	21,721	0.04	20,860
San Antonio, TX	1,547,253	2017–2020	92,444	1,026	0.01	2,460
Montgomery County, MD	1,050,688	2018–2020	170,313	12,225	0.07	6,808
San Jose, CA	1,021,795	2005–2020	273,199	10,364	0.04	16,026
San Francisco, CA	881,549	2017–2020	631,629	4,200	0.01	6,801
Seattle, WA	753,675	2010–2020	108,894	9,391	0.09	
Detroit, MI	670,031	2019–2020	929,216	21,125	0.02	8,209
Baltimore, MD	593,490	2014–2020	470,385	6,447	0.01	7,045
Tucson, AZ	548,073	2013–2020	193,581	20,585	0.11	10,770
Mesa, AZ	518,012	2017–2020	115,198	10,438	0.09	5,469
Sacramento, CA	513,624	2014–2020	298,380	9,068	0.03	
New Orleans, LA	390,144	2011–2020	398,042	16,496	0.04	10,932
Total (14 cities)	16,842,878		6,255,211	191,607	0.03	116,809

Notes: Neighborhoods is the total number of unique locations from which a domestic violence call originated. The geographic origin of each call is known up to the 100-block level. For example, “303 Main St” and “340 Main St” would both be recorded as originating from the same neighborhood (“300 Block of Main St”). The origin of domestic violence calls is not known for Los Angeles, Seattle, and Sacramento. Population comes from the 2019 U.S. Census estimates.

Out of the 52 million 911 calls, approximately 1.6 million (3%) are domestic violence-related. To determine whether a call was domestic violence-related, we tabulated the short textual descriptions that accompany each record, separately for each city, and manually read them to identify domestic violence-related descriptors. This process yielded 291 unique domestic violence-related descriptors across the 14 cities.⁸

The domestic violence calls come from 116,809 distinct neighborhoods. We constructed these neighborhoods from the call records by analyzing the “origin of call” information attached to each record. For 11 of the 14 cities, each call record includes information about the geographic origin of the call.⁹ The addresses in the call records are typically anonymized to the 100-level address block. For example, the locations “310 Chicago Ave,” “319 Chicago Ave,” and “375 Chicago Ave” would all be recorded as “300 Block Chicago Ave.” Some cities also use intersections to record locations (e.g. “Chicago Ave / State St”). We standardize these addresses to ensure consistency within a city and then use them to define unique neighborhood identifiers. Each neighborhood is thus about half of one city block or smaller.

Domestic violence calls exhibit strong seasonality (Figure 3). In general, call volumes increase linearly from the beginning of the calendar year, peak in early and late summer, and then linearly decline throughout the rest of the year. Four calendar days also stand out as outliers: July 4, July 5, and December 25 each exceed average domestic violence call levels by more than 20%, while January 1 exceeds the average level by more than 40%.

Importantly, however, the seasonal trend is stable throughout the period we study: It is roughly linear and continuous during the weeks before and after the calendar day of March 14 (which, for reasons explained below, is the day we use to mark the beginning of lockdown). In the regression analysis, we will use simple linear and quadratic trends to account for seasonality.

Domestic violence calls increased sharply in the weeks following lockdown. Figure 2 plots the average daily call volume in 2020. The average call

8. A list of these descriptors is given in the [Supplementary Appendix](#).

9. The origin of domestic violence calls is not known for Los Angeles, Seattle, and Sacramento.

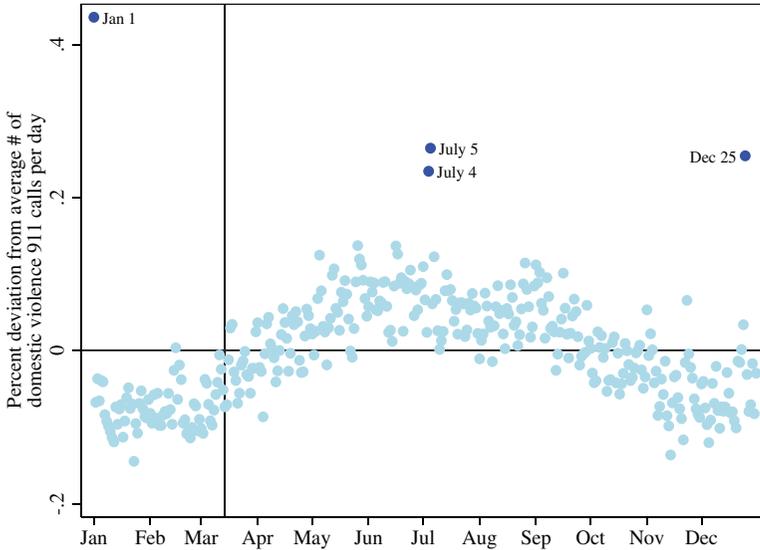


Figure 3. The seasonality of domestic violence 911 calls.

This figure shows that the seasonality of domestic violence 911 calls is, in years before lockdown, continuous throughout the pre- and post-lockdown period. Each dot is the average number of domestic violence 911 calls on that calendar day divided by the average number of domestic violence 911 calls (across all days). The sample is domestic violence 911 calls from all cities for all calendar days except February 29 and all years except 2020. The vertical line at March 14 indicates the start of the lockdown period in 2020.

volume follows the usual linear trend up to the week ending on March 13. It then sharply rises throughout lockdown and peaks in early April. By the end of April, the call volume declines to the level predicted by the pre-lockdown linear trend.

2.2. Mobile Device Location Data

The mobile device location data are a daily cross-sectional sample of mobile devices physically located within the 14 cities in our sample. The sample ranges from approximately 1.3 to 1.8 million devices per day, which is roughly one-tenth the size of the total population of the cities they represent. The data provider identifies a mobile device as being “at home” by comparing its current location to its history of nighttime locations for the

past 6 weeks.¹⁰ The mobile device location data are available beginning January 1, 2020.

We use the location data to estimate the total number of people at home for each hour of each day. A limitation of the location data is that it includes the number of devices found at home—but not the percent of devices that, when pinged, were found at home. Thus, to construct an estimate of the percent of devices at home, we must choose an appropriate denominator. We take the simple approach of dividing the total number of devices found at home in a given hour by the maximum number of devices found at home during that day. Since the maximum typically occurs during one of the early morning hours, this procedure normalizes the fraction at home to 1 for one of the early morning hours. Formally, our estimate of the total number of people at home at a given hour is

$$\text{Population} \times \frac{m_{ht}}{\max\{m_{h=1,t}, \dots, m_{h=24,t}\}}, \quad (1)$$

where m_{ht} is the number of mobile devices at home on hour h of date t .

Figure 1 plots our estimate of the fraction of people at home at 12 p.m. each day. The fraction at home on weekdays nearly doubled in mid-March, from about 45% to 85%. For weekends, the increase was smaller but similarly abrupt. The largest increase occurred on March 14. For this reason, we will use that day—March 14, 2020—to define the beginning of lockdown.

Since not everyone is at home during the early morning hours, our estimator of the fraction at home is biased upwards. To get a sense of the potential scale of this bias, suppose the maximum number of people at home occurs at 3 a.m. If the true fraction of people at home at 3 a.m. is 0.95, and the true fraction at home at 12 p.m. is 0.5, then our estimate at 3 a.m. would be biased upward by 5 percentage points ($= 95/95 - 95$) and our estimate at 12 p.m. would be biased upward by 3 percentage points ($\approx 50/0.95 - 50$).

More importantly, this procedure may also underestimate the change in the fraction at home caused by lockdown. Consider the change in at-home patterns on weekdays at 12 p.m. From Figure 1, the fraction at home increased 40 percentage points (from 0.45 to 0.85). If lockdown caused

10. The data were provided by SafeGraph. See the [Supplementary Appendix](#) for additional details.

the maximum fraction at home to increase from, say, 0.95 to 0.99, then the true change on weekdays at 12 p.m. would have been 41 percentage points (from 0.43 to 0.84). The bias for the weekday daytime change is therefore likely small. On the other hand, and continuing the example from above, suppose the maximum fraction at home always occurs at 3 a.m. Then the true change at 3 a.m. would be 4% (95–99), but our procedure would normalize this to zero by construction. In percentage terms, the scale of the downward bias is therefore modest for the main daytime estimates (where the percentage point changes are large), but potentially significant for the nighttime estimates (where the percentage point changes are small). We will keep this bias in mind when interpreting the results below.

3. The Impact of Lockdown on Domestic Violence

3.1. Baseline Results

We begin by estimating the total impact of lockdown on domestic violence. Our baseline specification is

$$\ln y_{ht} = \beta_0 + \beta_1 \text{Lockdown}_t + \beta_3 t + \beta_4 t^2 + \mu_{dow \times h} + u_{ht}, \quad (2)$$

where y_{ht} is the number of domestic violence 911 calls during hour h of date t , Lockdown_t is an indicator equal to 1 if t is on or after March 14, 2020, β_3 and β_4 are coefficients of a quadratic trend, and $\mu_{dow \times h}$ is a fixed effect for each cell defined by the interaction of day-of-week and hour-of-day ($7 \times 24 - 1$ total fixed effects). β_1 measures the percentage increase in domestic violence calls. The quadratic time trend is designed to control for the seasonality of 911 calls (Figure 3). To estimate equation 2, we use calls placed between January 2, 2020 to the end of the sample (April 24, 2020). We choose this time frame to ease comparison with later estimates that use the mobile device location data, which is available beginning January 1, 2020. We exclude January 1 because it is a consistent outlier.

Our baseline estimate of the increase in domestic violence caused by lockdown is 12% (Table 2, panel A). The estimate is similar (13%) even without controls for hour of day, day of week, or seasonal trends. Because lockdown had the greatest impact in at-home patterns during working hours,

we separately estimate equation 2 for periods defined by the interaction of daytime (9 a.m.–9 p.m.) and weekday (Monday–Friday). The largest increase occurred during daytime hours: 20% for weekdays and 17% for weekends. The nighttime increase was 9% for weekdays and less than 1% for weekends. The latter estimate is not statistically significantly different from zero at conventional levels of confidence.

3.2. The Effect of Increased Time Spent At Home

Since lockdown forced people to stay at home, it is natural to ask whether the impact of lockdown on domestic violence is, in percentage terms, equal to the impact of lockdown on at-home patterns (i.e. whether the elasticity of domestic violence with respect to people at home is equal to one). Further, it is also natural to ask whether the differential impacts we observe throughout the week are consistent with a constant-elasticity model.

To address these questions, we posit a standard two-stage process in which the log number of domestic violence 911 calls is a function of the number of people at home, and the log number of people at home is in turn a function of whether the community is under lockdown:

$$\ln y_{ht} = \gamma_0 + \gamma_1 \ln x_{ht} + u_{ht} \quad (3)$$

$$\ln x_{ht} = \pi_0 + \pi_1 \text{Lockdown} + v_{ht}, \quad (4)$$

where x_{ht} is the number of people at home during hour h on day t . Although this formulation is equivalent to using the lockdown period as an instrumental variable for time spent at home, it is worth emphasizing that we do *not* interpret γ_1 as the causal effect of being at home on domestic violence. To be sure, domestic violence typically occurs at home, and lockdowns exogenously forced people to stay home. But lockdowns were triggered by the coronavirus pandemic, which together with lockdown itself led to unemployment, anxiety, and myriad other forms of economic and psychological distress that directly impacted domestic violence. To be a valid instrument for time spent at home, lockdown would have had to impact domestic violence “only through” its impact on the number of people at home. Since this assumption is probably false, lockdown is not a valid instrument for the number of people at home.

Nevertheless, γ_1 has two related and useful interpretations. Firstly, and as a matter of arithmetic, γ_1 is the ratio of two causal effects: it is the reduced-form effect of lockdown on domestic violence (β_1 from equation 2) divided by the first-stage effect of lockdown on being at home (π_1 from equation 4). It therefore provides one way to scale the reduced-form estimates and compare the differential impacts of lockdown throughout the hours of the week. For example, if γ_1 were, say, 0.40, then we would conclude that the impact of lockdown on domestic violence was 40% of its impact on being at home.

The second and related interpretation of γ_1 is that, under plausible assumptions, it is an upper bound on the causal effect of being at home on domestic violence. To see this, recall that lockdown is not a valid instrument for being at home because it likely violates the “only through” assumption: lockdown caused more people to stay home, but it also caused economic and psychological distress that in turn may have directly affected rates of domestic violence. Yet suppose one partitions the impacts of lockdown on domestic violence into two channels: the first, which we will call the mechanical effect of lockdown, is the true effect of more people at home on domestic violence. The second, which we will call the collateral effects of lockdown, is the sum of all other nonmechanical influences of lockdown on domestic violence. The collateral effects include any channel—any economic or psychological stress—that increased the rate of domestic violence per person-hour at home. We refer to this set of effects as “collateral” because they were clearly not the intended effects of stay-at-home orders, private self-isolation, layoffs, and all the other public and private responses to the pandemic. If one assumes that the collateral effects jointly led to a net increase in domestic violence, then γ_1 is an upper bound on the mechanical effect because it attributes the entire change in domestic violence to the change in time spent at home. Put another way, γ_1 is the mechanical effect under the (invalid) assumption that lockdown increased domestic violence “only through” its increase on the number of people at home.

Moving to the estimates, panel B of Table 2 reports results from the first stage: the effect of lockdown on the number of people at home. As one might expect, the largest increase in at-home patterns occurred during regular working hours. The number of people at home for all hours increased 14 percent. The increase during daytime, however, was 32% for weekdays

and 13% for weekends. The increases during nighttime were much smaller: 5% for weekdays and less than 1% in magnitude for weekends. It is worth noting, however, that the nighttime weekend estimate includes the hour-of-day and quadratic time trend. Without these controls, the simple post-lockdown difference is 2% and statistically significant.

For nighttime hours, there are two reasons to prefer the simple difference estimates without controls. First, as explained in Section 2, the upward bias inherent in our estimate of the number of people at home, while small in terms of percentage *point* differences, increases in percentage differences as the fraction of people at home goes to 1. This could lead us to underestimate the change in at-home patterns, especially for the early morning hours. On the other hand, and partially mitigating the magnitude of this potentially large bias, the change in at-home patterns is substantial during many of the hours that are included in the nighttime estimates, specifically, during the hours between 7 a.m.–9 a.m. and 9 p.m.–11 p.m. Second, the change in at-home patterns at night are likely small to begin with, as most people are probably asleep during at least half the hours between 9 p.m. and 9 a.m. The extrapolation of any pre-lockdown trend could therefore overwhelm these small differences, leading to a zero (or even negative) estimate (Bound et al., 1995). For these reasons, we interpret the nighttime first- and second-stage estimates with caution, as the first-stage estimate may be biased downward. At the same time, however, one may be more confident in interpreting the daytime estimates because the impact of the two concerns outlined above is likely to be much smaller.

Moving to the second-stage estimates (Table 2, panel C), the upper bound on the elasticity of domestic violence with respect to people at home is 0.85 (column 2). Thus, in percentage terms, the total change in domestic violence was 85% of the total change in at-home patterns. We fail to reject the null hypothesis that the elasticity is equal to 1, using both a basic *t*-test and the more accurate Anderson and Rubin (1949) test.¹¹

11. The *p*-value for the Anderson and Rubin (1949) test is $p = 0.29$ (column 2). Generally for columns 1 through 6, the first-stage relationship is strong enough that there is little difference between an asymptotic *t*-test and an Anderson and Rubin (1949) test. However, for the results without controls (column 1), we reject the null that the elasticity is equal to 1 using either approach.

With full controls, the daytime estimates of γ_1 are 0.62 for weekdays and 1.29 for weekends. Relative to the change in at-home patterns, the increase in domestic violence during weekday daytime hours, while larger in absolute terms, is thus smaller than the increase during weekend daytime hours. The difference in these estimates is statistically significant at conventional levels ($p = 0.02$, column 4 versus column 6).¹²

In summary, the results provide evidence against a simple proportional model. While we cannot reject the simple proportional model on average (column 2), we can reject equality of the proportions during daytime hours for weekdays versus weekends (columns 4 versus 6). We conclude that the lockdown-associated surge in domestic violence likely reflects a mixture of causal channels more complex than the simple mechanical increase in the number of people at home.

4. The Impact of Lockdown on the Extensive versus Intensive Margin of Domestic Violence

Finally, we ask whether the surge in domestic violence came from households with or without a history of domestic violence. To answer this question, one could use a panel of individual- or household-level data to estimate the impact of lockdown on the extensive versus intensive margin. The change in households without prior domestic violence calls would be the change in the extensive margin, while the change in households with prior domestic violence calls would be the change in the intensive margin. Unfortunately, the 911 call data do not include an individual or household identifier. They do include, however, addresses that were coarsened to the level of about one half of one city block, which we use to construct neighborhood-level identifiers. In this section, we show how the neighborhood-level data can provide an approximation of the change in the extensive margin of domestic violence, as well as an upper bound for the change in the intensive margin.

12. However, without controls, the estimates are much more similar (0.52 for weekdays, 0.61 for weekends). Given the sampling variability of each estimate, it is perhaps not surprising that the data suggest little difference between the two estimates ($p = 0.60$, column 3 versus column 5).

4.1. Theoretical framework

We begin by defining the extensive and intensive margins of domestic violence. We will say that a call occurs on the extensive margin if it is the first call for that household since period $t = 0$, for some start date normalized to zero. Otherwise, the call occurs on the intensive margin.

It is worth pausing here to clarify the role that the “start date” plays in both the theoretical model and the empirical application. In principle, the first “true” extensive margin event for a given household is the first time that the household places a 911 call. Given household-level data, the ideal start date would therefore be the date that the household was formed. However, this is not feasible in our setting because the call data are at the neighborhood level. It is for this reason that we redefine the first extensive margin event as the first call *since a given start date*, and then show (in Appendix Section A) how changes in the (redefined) neighborhood margins can identify changes in the (redefined) household margins. Thus, the interpretation of the results depends on the specific choice of start date. Since the exact choice of start date is arbitrary, the empirical section repeats the analysis for a variety of start dates.

Returning to the model, let θ_i denote the probability that household i places a domestic violence 911 call in a given period. The expected number of calls in period t is then

$$\mu_t = \sum_i \theta_i. \quad (5)$$

The expected contribution of household i to the extensive margin is

$$\mu_{it}^{ext} = \theta_i \cdot \prod_{\tau=1}^t (1 - \theta_i) \quad (6)$$

and the expected contribution of household i to the intensive margin is

$$\mu_{it}^{int} = \theta_i \cdot \left[1 - \prod_{\tau=1}^t (1 - \theta_i) \right]. \quad (7)$$

We model lockdown as a shock that occurs in period s and scales up θ_i for all $t \geq s$. The scaling factor is allowed to differ for domestic violence calls on the extensive versus intensive margin. Let η denote the scaling factor for

the extensive margin, and ε the scaling factor for the intensive margin. After the shock, household i 's expected contribution to the extensive margin is

$$\mu_{it \geq s}^{ext} = (1 - \theta_i)^s (1 - \eta \cdot \theta_i)^{t-s} \cdot \theta_i \cdot \eta \quad (8)$$

and its expected contribution to the intensive margin is

$$\mu_{it \geq s}^{int} = [1 - (1 - \theta_i)^s (1 - \eta \cdot \theta_i)^{t-s}] \cdot \theta_i \cdot \varepsilon. \quad (9)$$

This motivates the following econometric modeling of the log aggregate extensive and intensive margin counts in period t :

$$\ln y_t^{ext} \approx \kappa_1 + I(t \geq s) \cdot \ln \eta + \kappa_2 \cdot t \quad (10)$$

and

$$\ln y_t^{int} \approx \kappa_3 + I(t \geq s) \cdot \ln \varepsilon, \quad (11)$$

where y_t^{ext} is the total number of calls from households without a call since period $t = 0$, y_t^{int} is the total number of calls from households with at least one call since period $t = 0$, $I(\cdot)$ is the indicator function, and κ_1 , κ_2 , and κ_3 are constants. After the shock, the magnitude of each margin shifts by an amount approximately equal to (one minus) the scaling factors η and ε .

As noted, a limitation of our data is that it identifies calls at a neighborhood level, but not a household level. However, neighborhood level data contain valuable information for the intensive–extensive breakdown of our results. An increase in calls from a neighborhood with no history of calls can only occur on the extensive margin. An increase in calls from a neighborhood with a history of calls, on the other hand, could come from either the intensive or extensive margin, and therefore represents a mixture of the two. Appendix Section A develops the implications of this intuition. Letting \sim denote neighborhood-level quantities, Appendix Section A shows that $\tilde{\eta}$ is a reasonable approximation of η and that $\tilde{\varepsilon}$ is a bound (upper or lower, depending on η) on ε .

In practice, it may be that the impact of lockdown on the extensive margin differs by neighborhood. In that case, our estimate of η is valid for the households in neighborhoods with low rates of domestic violence. Note that a neighborhood-level rate can be low either because the rates of its

constituent households are low, or because there are simply few households in the neighborhood.

4.2. Results

To estimate the impact of lockdown on the extensive and intensive margins, we use the specifications implied by the theoretical framework:¹³

$$\ln y_t^{ext} = \phi_0 + \phi_1 \text{Lockdown}_t + \phi_2 t + u_t \quad (12)$$

$$\ln y_t^{int} = \psi_0 + \psi_1 \text{Lockdown}_t + \psi_2 t + v_t. \quad (13)$$

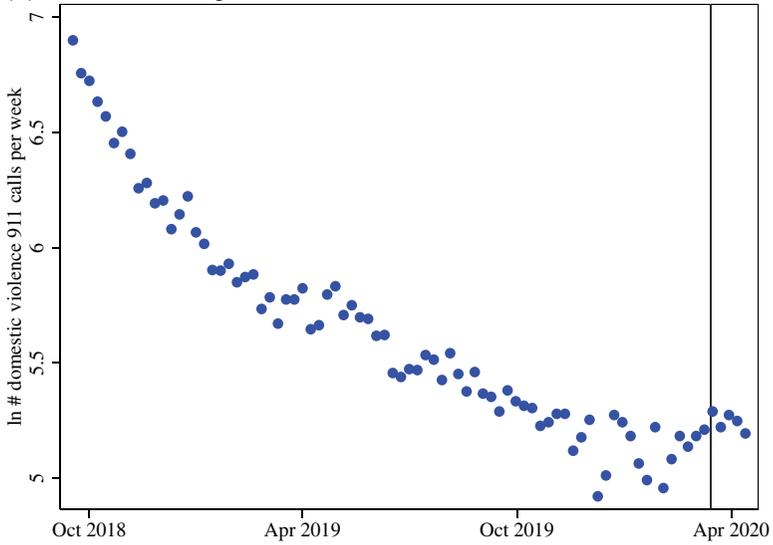
The change in the extensive margin is ϕ_1 , and the bound on the change in the intensive margin is ψ_1 .

As a baseline, we use 18 months before the lockdown period (September 14, 2018) as the start date to define the extensive and intensive margins. A domestic violence 911 call thus contributes to the extensive margin if it is the first neighborhood-level call since September 14, 2018. Otherwise, it contributes to the intensive margin. In addition, since the functional form approximations from equations 10 and 11 hold only when t is sufficiently large relative to θ (when one is sufficiently far from the start date that defines both margins), we only use observations beginning 150 days after the start date to estimate both equations. The approximations appear to be valid about 150 days after the start date, at which point the extensive margin is roughly linear and the intensive margin is roughly constant (figure 4). Finally, since the choice of start date and estimation window is arbitrary, we also produce estimates for other start dates and estimation windows.

The impact on the extensive margin is consistently larger than the (upper) bound of the impact on the intensive margin. Table 3 reports estimates using the start date September 14, 2018 and commencing estimation 150 days later. The change in the extensive margin is larger on average (16 percent versus 8 percent) as well as for each of the sub-periods during the week. During daytime, the increase was 23 versus 14 percent for weekdays, and

13. In our simple model, $\psi_2 \approx 0$ for sufficiently large t . See section A. Also, the estimates below do not include Detroit because data are not available before January 1, 2019. Similar results are obtained when one restricts the start dates to January 1, 2019 and later. (Results not shown.)

(a) Extensive margin



(b) Intensive margin

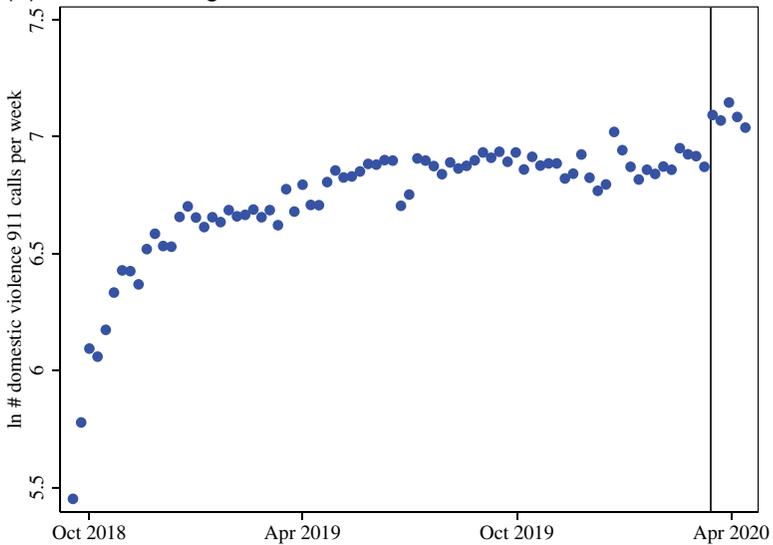


Figure 4. The impact of lockdown on (a) the extensive margin and (b) the intensive margin of domestic violence (Monday–Friday, 9 a.m.–9 p.m.). A call is on the extensive margin if it is the first call from that neighborhood since September 14, 2018 (18 months before lockdown begins). Otherwise, the call is on the intensive margin.

Table 3. The Impact of Lockdown on Domestic Violence (Extensive versus Intensive Margins)

	Dependent variable is log number of domestic violence-related 911 calls on a given day and hour											
	All hours				Daytime (9 a.m.–9 p.m.)				Nighttime (9 p.m.–9 a.m.)			
	Ext.	Int.	Ext.	Int.	Ext.	Int.	Ext.	Int.	Ext.	Int.	Ext.	Int.
Lockdown	0.16** (0.02)	0.08** (0.01)	0.23** (0.03)	0.14** (0.02)	0.25** (0.04)	0.10** (0.02)	0.12** (0.03)	0.06** (0.02)	0.01 (0.05)	0.01 (0.05)	0.01 (0.05)	-0.01 (0.03)
Hour × Day of week f.e.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R ²	0.44	0.63	0.27	0.41	0.27	0.24	0.45	0.64	0.46	0.46	0.46	0.64
Observations	10,511	10,511	3,768	3,768	1,488	1,488	3,767	3,767	1,488	1,488	1,488	1,488

Notes. This table estimates the impact of lockdown on the log number of domestic violence 911 calls for the extensive and intensive margins. A call is on the extensive margin if it is the first call from that neighborhood since September 14, 2018 (18 months before lockdown begins). Otherwise, the call is on the intensive margin. Lockdown begins March 14, 2020. * and ** indicate statistically significantly different from zero at 95% and 99% confidence, respectively.

25 versus 10 percent for weekends. During nighttime, the changes were smaller: 12 versus 6 percent for weekdays, and 1 versus -1 for weekends. (The latter result is not statistically significant.) Given that our estimate of the intensive margin change is an upper bound, these results suggest that the percentage change in the extensive margin is roughly twice as large (or more) than the change in the intensive margin.

Finally, the main result – that the impact on the extensive margin was greater than the impact on the intensive margin – is robust to a variety of start dates and estimation windows. Specifically, it is robust to using start dates between 12 and 24 months before lockdown, and estimation windows that begin 120, 150, and 180 days after the start date. In theory, one might prefer the earlier start dates as they are more likely to capture a household’s true first domestic violence call (and thus reflect the true change in the extensive margin). On the other hand, an earlier start date allows more time for household and neighborhood turnover to bias the estimates in unknown directions. In table 3, we chose 18 months before lockdown as an (admittedly arbitrary) balance between these two concerns. Somewhat reassuringly, however, estimates of changes on the extensive margin are fairly stable, particularly for start dates that are 18 months or more before lockdown.¹⁴

5. Conclusion

We used 911 calls for police service to study the impact of the coronavirus lockdown on domestic violence. Domestic violence calls surged in the weeks following lockdown, from mid-March to early-April, and returned to pre-lockdown levels by the end of April. The largest increases occurred during weekday daytime hours, when most adults and children would have otherwise been at work or in school. Using mobile device location data, we found that these hours also experienced the greatest disruption in at-home patterns. Lockdown is especially likely to have led to episodes of first-time abuse: the increase from neighborhoods with no recent history of violence was roughly double the increase from neighborhoods with a recent history of violence.

14. See the [Supplementary Appendix](#).

Our results are important for evaluating the total social cost of prolonged lockdown policies. Compared to the economic costs of lost GDP and unemployment, the physical and psychological costs of lockdown are much more difficult to observe and quantify. They are, however, just as important contributors to welfare.

A. Technical Appendix

This appendix section formally shows how we use neighborhood-level data to identify household-level changes on the extensive and intensive margins.

Household-level margins. For plausible baseline call rates,¹⁵

$$(1 - \eta \cdot \theta_i)^{t-s} \approx (1 - \theta_i)^{t-s} \quad (\text{A.1})$$

for t close to s , and

$$(1 - \theta_i)^t \approx 0 \quad (\text{A.2})$$

for sufficiently large t . Thus

$$\ln y_t^{ext} \approx \ln \sum_i (1 - \theta_i)^t \cdot \theta_i + I(t \geq s) \cdot \ln \eta \quad (\text{A.3})$$

$$\ln y_t^{int} \approx \ln \sum_i \theta_i + I(t \geq s) \cdot \ln \varepsilon. \quad (\text{A.4})$$

Neighborhood-level margins. We observe the neighborhood-level extensive and intensive margins:

$$\tilde{\mu}_t^{ext} = \begin{cases} \sum_g (1 - \theta_g)^t \cdot \theta_g & \text{for } t < s \\ \sum_g (1 - \theta_g)^s (1 - \tilde{\eta} \cdot \theta_g)^{t-s} \cdot \theta_i \cdot \tilde{\eta} & \text{for } t \geq s \end{cases} \quad (\text{A.5})$$

and

$$\tilde{\mu}_t^{int} = \begin{cases} \sum_g [1 - (1 - \theta_g)^t] \cdot \theta_g & \text{for } t < s \\ \sum_g [1 - (1 - \theta_g)^s (1 - \tilde{\eta} \cdot \theta_g)^{t-s}] \cdot \theta_i \cdot \tilde{\varepsilon} & \text{for } t \geq s, \end{cases} \quad (\text{A.6})$$

15. The average daily call rate is about 3–5 per 100,000.

where

$$1 - \theta_g = \prod_{i \in g} (1 - \theta_i) \tag{A.7}$$

is the probability of zero calls in neighborhood g in a given time period.

Consider the extensive margin scaling factor for neighborhood g , denoted $\tilde{\eta}_g$. Suppose there are n persons in neighborhood g , and make the simplifying assumption that $\theta_i = \theta \forall i \in g$.¹⁶ Then the percentage change in the neighborhood-level extensive margin is

$$\begin{aligned} \tilde{\eta} &= \frac{1 - \prod_{i \in g} (1 - \eta \cdot \theta_i)}{1 - \prod_{i \in g} (1 - \theta_i)} \\ &= \frac{1 - (1 - \eta \cdot \theta)^n}{1 - (1 - \theta)^n} \\ &= \eta \times \frac{\sum_{k=1}^n \binom{n}{k} \theta^k (-\eta)^{k-1}}{\sum_{k=1}^n \binom{n}{k} \theta^k (-1)^{k-1}} \\ &= \eta \times \frac{\binom{n}{1} \theta - \binom{n}{2} \theta^2 \eta + \binom{n}{3} \theta^3 \eta^2 - \binom{n}{4} \theta^4 \eta^3 + \dots}{\binom{n}{1} \theta - \binom{n}{2} \theta^2 + \binom{n}{3} \theta^3 - \binom{n}{4} \theta^4 + \dots} \\ &\approx \eta. \end{aligned} \tag{A.8}$$

As an estimate of η , $\tilde{\eta}$ has a second-order downward bias, a third-order upward bias, a fourth-order downward bias, and so on. The net bias is downward, and the magnitude increases with η , n , and θ . For plausible values of η , n , and θ , however, the bias is negligible. The order of magnitude of the bias is roughly $\theta \cdot n$. In neighborhoods without recent calls, it is likely that the typical household call rate is well below average. Nevertheless, using the conservatively high level of the average daily call rate from the data (about 0.00004) and a neighborhood size of 25, a household-level η of 1.5 would yield a neighborhood-level $\tilde{\eta}$ of approximately 1.4996. A neighborhood of 250 would be 1.4963.

Next consider the neighborhood-level intensive margin. In pre-shock periods we observe

$$\tilde{\mu}_{t < s}^{int} = \sum_g (1 - (1 - \theta_g)^t) \cdot \theta_g \tag{A.9}$$

16. This is without loss of generality given sufficiently small θ .

and in the period of the shock we observe

$$\tilde{\mu}_{t=s}^{int} = \sum_g (1 - (1 - \theta_g)^t) \cdot \theta_g \cdot \tilde{\varepsilon}. \quad (\text{A.10})$$

The (immediate) percentage change in the intensive margin for neighborhood g is

$$\begin{aligned} \tilde{\varepsilon}_g &= \frac{1 - \prod_{i \in g} \{(1 - \theta_i)^s (1 - \eta \cdot \theta_i) + (1 - (1 - \theta_i)^s)(1 - \varepsilon \cdot \theta_i)\}}{1 - \prod_{i \in g} (1 - \theta_i)} \\ &= \frac{1 - \prod_{i \in g} \{1 - (\varepsilon + a \cdot (\eta - \varepsilon)) \cdot \theta_i\}}{1 - \prod_{i \in g} (1 - \theta_i)}, \end{aligned} \quad (\text{A.11})$$

where $a = (1 - \theta_i)^s$. The neighborhood-level intensive margin is thus biased toward η .

$\tilde{\varepsilon}$ is a reasonable approximation of ε only in three special cases: (1) when all persons in neighborhood g have had a recent call ($a = 0$), (2) when persons in g without a recent call have baseline rates of zero ($\theta_i|_{a>0} = 0$), and (3) the knife-edge case when $\varepsilon = \eta$. In cases (1) and (2), post-shock calls in neighborhoods with a recent call can be attributed to the intensive margin. In case (3), the neighborhood-level elasticity is “biased” toward the household-level elasticity. In the general case when neither (1) nor (2) nor (3) hold, a comparison of $\tilde{\varepsilon}$ and $\tilde{\eta}$ will determine whether $\tilde{\varepsilon}$ is an upper or lower bound on ε . If $\tilde{\varepsilon} > \eta$ then

$$\eta < \tilde{\varepsilon} < \varepsilon \quad (\text{A.12})$$

and $\tilde{\varepsilon}$ is a lower bound on ε , while if $\tilde{\varepsilon} < \eta$ then

$$\varepsilon < \tilde{\varepsilon} < \eta \quad (\text{A.13})$$

and $\tilde{\varepsilon}$ is an upper bound on ε .

Supplementary material

Supplementary material is available at *American Law and Economics Review* Journal online.

References

- Anderson, Theodore W., and Herman Rubin. 1949. “Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations,” 20 *The Annals of Mathematical Statistics* 46–63.
- Balsam, Kimberly F., and Dawn M. Szymanski. 2005. “Relationship Quality and Domestic Violence in Women’s Same-Sex Relationships: The Role of Minority Stress,” 29 *Psychology of Women Quarterly* 258–269.
- Boserup, Brad, Mark McKenney, and Adel Elkbuli. 2020. “Alarming Trends in US Domestic Violence during the COVID-19 Pandemic,” 38 *The American Journal of Emergency Medicine* 2753–2755.
- Bound, John, David A. Jaeger, and Regina M. Baker. 1995. “Problems With Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogenous Explanatory Variable is Weak,” 90 *Journal of the American Statistical Association* 443–450.
- Calderon-Anyosa, Renzo J. C., and Jay S. Kaufman. 2020. “Impact of COVID-19 Lockdown Policy on Homicide, Suicide, and Motor Vehicle Deaths in Peru,” 143 *Preventive Medicine* 106331.
- Dahal, Minakshi, Pratik Khanal, Sajana Maharjan, Bindu Panthi, and Sushil Nepal. 2020. “Mitigating Violence against Women and Young Girls during COVID-19 Induced Lockdown in Nepal: A Wake-up Call,” 16 *Globalization and Health* 1–3.
- Das, Manob, Arijit Das, and Ashis Mandal. 2020. “Examining the Impact of Lockdown (Due to COVID-19) on Domestic Violence (DV): An Evidences from India,” 54 *Asian Journal of Psychiatry* 102335.
- Edleson, Jeffrey L. 1999. “Children’s Witnessing of Adult Domestic Violence,” 14 *Journal of Interpersonal Violence* 839–870.
- Gillum, Jack, Lisa Song, and Jeff Kao. 2020. “There’s Been a Spike in People Dying at Home in Several Cities. That Suggests Coronavirus Deaths Are Higher Than Reported,” *ProPublica* April 14, 2020.
- Golding, Jacqueline M. 1999. “Intimate Partner Violence as a Risk Factor for Mental Disorders: A Meta-Analysis,” 14 *Journal of Family Violence* 99–132.
- Hanmer, Jalna, Sue Griffiths, and David Jerwood. 1999. *Arresting Evidence: Domestic Violence and Repeat Victimization*, Vol. 104. Home Office London.
- Heckman, James J. 1981. “Heterogeneity and State Dependence” in Sherwin Rosen, ed., *Studies in Labor Markets*. Chicago, IL: University of Chicago Press, pp. 91–140.
- Javed, Saira, and Yasir Mehmood. 2020. “No Lockdown for Domestic Violence during COVID-19: A Systematic Review for the Implication of Mental-Well Being,” 1 *Life and Science* 6.

- Kilpatrick, Kym L., and Leanne M. Williams. 1997. "Post-traumatic Stress Disorder in Child Witnesses to Domestic Violence," 67 *American Journal of Orthopsychiatry* 639.
- Leslie, Emily, and Riley Wilson. 2020. "Sheltering in Place and Domestic Violence: Evidence from Calls for Service during COVID-19," 189 *Journal of Public Economics* 104241.
- Ott, Haley. 2020. "In Countries Mired in Crises, Aid Organization Sees Huge Drop in Domestic Abuse Reports – and It's Worried," *CBS News*.
- Perez-Vincent, Santiago M., Enrique Carreras, M. Amelia Gibbons, Tommy E. Murphy, and Martín A. Rossi. 2020. *COVID-19 Lockdowns and Domestic Violence*. Washington, DC, USA: Inter-American Development Bank.
- Ravindran, Saravana, and Manisha Shah. 2020. "Unintended Consequences of Lockdowns: Covid-19 and the Shadow Pandemic," Technical Report, National Bureau of Economic Research.
- Russell, Diana E. H. 1990. *Rape in Marriage*. Bloomington, IN: Indiana University Press.
- Sediri, Sabrine, Yosra Zgueb, Sami Ouanes, Uta Ouali, Soumaya Bourgo, Rabaa Jomli, and Fethi Nacef, 2020. "Women's Mental Health: Acute Impact of COVID-19 Pandemic on Domestic Violence," 23 *Archives of Women's Mental Health* 749–756.
- Sifat, Ridwan Islam. 2020. "Impact of the COVID-19 Pandemic on Domestic Violence in Bangladesh," 53 *Asian Journal of Psychiatry* 102393.
- Snyder, James, Abigail Gewirtz, Lynn Schrepferman, Suzanne R Gird, Jamie Quattlebaum, Michael R. Pauldine, Katie Elish, Osnat Zamir, and Charles Hayes. 2016. "Parent–Child Relationship Quality and Family Transmission of Parent Posttraumatic Stress Disorder Symptoms and Child Externalizing and Internalizing Symptoms Following Fathers' Exposure to Combat Trauma," 28 *Development and Psychopathology* 947–969.
- Stevenson, Betsey, and Justin Wolfers. 2009. "The Paradox of Declining Female Happiness," 1 *American Economic Journal: Economic Policy* 190–225.
- Truman, Jennifer L., and Rachel E. Morgan. 2014. "Nonfatal Domestic Violence, 2003–2012," Technical Report, Washington, D.C.
- White, Jacquelyn W., and Mary P. Koss. 1991. "Courtship Violence: Incidence in a National Sample of Higher Education Students," 6 *Violence and Victims*, 247.