NBER WORKING PAPER SERIES

DETERRENCE THROUGH RESPONSE CURVES:
AN EMPIRICAL ANALYSIS OF THE GAZA-ISRAEL CONFLICT

Eli Berman
Prabin B. Khadka
Danny Klinenberg
Esteban Klor

Working Paper 33273
http://www.nber.org/papers/w33273

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2024

We acknowledge the support of the Minerva Research Institute (FA9550-23-1-0437). The views expressed in this article are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. Esteban F. Klor thanks the Israel Science Foundation (grant No. 1061/18) and the Falk Institute for Economic Research in Israel for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Deterrence Through Response Curves: An Empirical Analysis of the Gaza-Israel Conflict
Eli Berman, Prabin B. Khadka, Danny Klinenberg, and Esteban Klor
NBER Working Paper No. 33273
December 2024
JEL No. C73, F50

ABSTRACT

We use response curves in a repeated game to analyze key aspects of mutual deterrence: escalation, de-escalation, incomplete deterrence, and deterrence by denial. In this approach, episodes of violence are due to interacting response curves, which disincentivize opponents from attacking through both deterrence and compellence. Both sides punish attacks to maintain credibility in future episodes, disincentivizing larger attacks and yielding nonviolent lulls. We empirically estimate those curves using detailed incident data from the Israel-Gaza conflict between 2007 and 2017. Our estimates match the dynamics of the raw data: very frequent episodes of low lethality violent exchange. Response curves are stable and probably Markovian. They exhibit a posture consistent with incomplete deterrence: i.e., episodes de-escalate, but to a violent equilibrium. Israeli missile defense shifts the Gazan response curve to a less aggressive posture, as predicted by theory.

Eli Berman
Department of Economics, 508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093
and NBER
elib@ucsd.edu

Prabin B. Khadka
University of Essex
prabin.khadka@essex.ac.uk

Danny Klinenberg
University of California, San Diego
dklinenberg@ucsb.edu

Esteban Klor
Dept. of Economics Hebrew
University Mount Scopus
Jerusalem ISRAEL 91905
eklor@huji.ac.il

1. Introduction

Deterring adversaries is fundamental to security. Yet deterrence is often incomplete: even the most powerful actors suffer a plethora of low-level attacks and provocations, from espionage, cyber-attacks, and election interference to proxy terrorism and missile tests. We use the analytical device of response curves in a repeated game to examine key aspects of deterrence relevant to modern conflict: escalation, de-escalation, incomplete deterrence, and deterrence by denial. Response curves (equivalently, reaction curves) draw on a rich literature in noncooperative game theory (Schelling 1960 and 1966; Fudenberg and Tirole 1991).[1]

Our approach allows us to empirically estimate response curves and to test theoretical predictions about deterrence by denial using high-frequency data from the Israel-Gaza conflict, covering the period 2007-2017. That period contains long intervals of relatively low-intensity violence at high frequency (averaging an episode of violent exchange every week) interrupted by three large Israeli incursions into Gaza, aimed at destroying Hamas' offensive capacity.

In our model, episodes of violence are due to the interaction of stable response curves. Response curves follow a strategic logic of disincentivizing opponents from attacking through deterrence and compellence. To maintain credibility, both sides punish attacks, which result in periods of relative quiet, and disincentivize larger attacks. Deterrence can be incomplete (i.e., violence occurs in equilibrium) when at least one adversary faces high net costs of suppressing low levels of violence, while another has positive net returns to low-level violence. We don't take a strong stand on the underlying social welfare functions from which response curves might be derived. A response curve approach to analyzing episodes of conflict can encompass both rationalist and other explanations for conflict, such as leaders appeasing domestic audiences who might seek retribution (Fearon 1994), or conflict due to psychological biases or over-responses of combatants and decision-makers (Jervis, Lebow and Stein 1985).[2]

Much game-theoretic work focuses on a single crisis where the threat of conflict or escalation is greatly influenced by some unobserved characteristic of an opponent, which shapes

---

[1] The idea of partial deterrence is developed in Freedman (2004).

[2] Building on the original insights of Schelling (1960, 1966), Kahn (1960), Mearsheimer (1983) and Powell (1990), several formal game-theoretic models of deterrence have been proposed to tie up loose ends regarding complexities of credibility. Subsequent studies have yielded additional significant theoretical breakthroughs, providing rigorous game-theoretic foundations for analysis of current policy questions such as how to promote state building, or whether and how to militarily intervene to induce regime change (Fearon 1997 and 2002; Powell 1999, 2003, 2012; Myerson 2008; Baliga and Sjöström 2009 and Chassang and Padró i Miquel 2010).

their response curve through perceived costs and benefits of attack. The unobserved characteristic could be the resolve of the opponent (i.e., the importance that the opponent gives to a particular issue under bargaining) or her military ability, which would in turn affect her (unknown) propensity to escalate attacks. This approach seems less compelling in long-standing conflicts such as Israel-Gaza, where opponents have plenty of opportunity to learn each other's characteristics.

The period of incomplete deterrence that we study came to a stunning end on October 7, 2023, so some context is needed before turning to results. First, as the literature on strategic surprise points out, deterrence is achieved when an adversary perceives that costs of a violent attack exceed benefits. Yet that perception, in the mind of an adversary, might well be hidden. Specifically, an adversary might achieve surprise (once deterrence has failed), by masking their true perception of costs and benefits until they attack (e.g., Chassang and Padro i Miquel, 2010). That seems to have been what Hamas leadership in Gaza achieved during the two years before the October 7th attacks.[3] Second, our data describe a period including Yahya Sinwar's election to lead Hamas in early 2017, but preceding documented preparations for the October 7, 2023 terrorist attack, which apparently began sometime in 2021 (though a similar plan existed already in 2014).[4] During our period of analysis (2007-2017) our working assumption is that Hamas' response curves were not a façade, but instead reflected a stable underlying objective function.[5]

We draw original data from daily reports as recorded by the United Nations Field Security Office. The reports describe security incidents in one or two sentences of text, followed by a description of any damage or casualties. Each description is preceded by a time stamp that is accurate to 5 minutes and is organized into four geographical areas. We code those data into "actions" (multiple attacks using air strikes, small arms fire, shelling, and incursions by the IDF,

[3] We lack corroborated evidence on why Hamas' leadership in Gaza (principally Yahya Sinwar and Mohammed Deif) perceived that the costs of a large attack on civilians no longer exceeded benefits by October 2023. Speculation based on comments by Hamas' leadership abroad include: (i) increased fear of a peace agreement between Israel and Saudi Arabia that would favor their political rivals in the Palestinian Authority, and doom Hamas' stated goal of liberating Palestine; and (ii) a calculation that a very predictable toll in civilian casualties and casualties among their own fighters, massive as it might be, would be justified by progress in the overarching objective of liberating the holy places (Hubbard and Abi-Habib, 2023). They may have also believed that Hamas' persistent accumulation of an infrastructure of tunnels, fortifications, and mobilized fighters would deter the IDF from a massive military incursion, or protect them from such an attack.

[4] Bergman and Goldman, 2023; Hila'i, 2023.

[5] We describe below why we choose to end the sample period in mid-December 2017, out of concern that Hamas may have been presenting a façade while planning the October 7, 2023 attacks.

and using small arms fire, rocket fire, and mortar fire by Gazan militants). We organize those actions sequentially into episodes of attack and counterattacks, separated by lulls. (Coding procedures are informed by social media of perpetrators.) We calculate predicted damage suffered by each side from a given attack, by associating each type of attack with actual casualties.

That approach yields a very high number of relatively low casualty episodes: 1,170 over our eleven years of data. The mean episode lasts 0.96 days. It includes 1.56 Gazan actions and 1.82 Israeli, resulting in an average of 0. 735 fatalities. We observe an average of only 12 days of lull between violent episodes.

Using those data we estimate response curves, which map expected damage suffered into expected damage caused. Both sides' response is consistent with mixed strategies, in the sense that they do not always respond to an attack, but their probability of response increases in expected damage suffered. Expected response curves are upwards sloping, indicating a deterrent posture; i.e., expected damage caused increases in damage suffered. Importantly, they are concave, so that they yield a single stable equilibrium. And their slopes are shallow enough that episodes de-escalate from high violence attacks back to equilibrium. Unfortunately, that equilibrium is at positive damage for both sides, a consequence of both intercepts being far from the origin. That is, our estimates exhibit incomplete deterrence in equilibrium. In particular, zero violence is not an equilibrium.

The introduction of Iron Dome missile defense by Israel in April 2011 shifts the Gazan response curve to lower levels of violence, illustrating deterrence by denial –i.e., by hardening targets, and yielding a less violent equilibrium. We lack sufficient precision to estimate whether Israel's major ground operations do the same.

Studies of conflict dynamics (see, for example, Powell 2012 or Leventoğlu and Slantchev 2007) are typically informed by a collection of cases. Our ability to estimate response curves using repeated interactions between the same opponents provides a unique empirical and analytical opportunity to illustrate both dynamics and equilibria in a stable system.

In the economic literature, the natural analogy to our estimation approach would be a sequence of pricing decisions by duopolies in a repeated game (e.g., Bajari et al., 2007). In Political Science or Sociology, this approach follows the "sequential analysis" tradition (Lichbach 1987, Heise, 1989). When applied to two-sided conflict, this approach tends to use a tit-for-tat repeated game as a model (Axelrod, 1984) and codes observational data to fit the structure of alternating sequential play in a repeated game. Moore (1998) provides an example from the literature on the repression of violent protests. Our innovations, enabled by data quality, are in modeling both probability of attack and attack intensity, and in estimating equilibria.

One novel contribution of our study is to test two key assumptions of the repeated game approach: that response curves are Markovian and, therefore, stable both within and between episodes. That test requires some nontrivial econometric analysis of serial correlation in sequences of actions and episodes, which we explain below. While those assumptions are technically rejected between episodes, we find that response functions are stable (i.e., Markovian) enough to provide an excellent approximation in studying mutual deterrence.

These results differ from those in the empirical literature on the Israel-Palestine conflict. Where we find de-escalation within a few days on average, that literature finds responses to an initial provocation that last over a month. Seminal empirical studies of the Israel-Palestine conflict have studied Israeli and Palestinian responses to the other side's attacks (Jaeger and Paserman 2006, 2008; Haushofer et al, 2010). Yet these efforts lacked frequent enough measurement to estimate response curves –they had available at best daily data, and our average episode lasts just less than a day. Those studies used principally Vector Autoregression (VAR). Haushofer et al. (2010) extended the Jaeger and Paserman (2008) analysis to include non-lethal acts of retaliation (e.g., Qassam rocket fire) using data from Gaza as well as from the West Bank. They find that optimal lag length varies for each outcome. They conclude that Israeli military actions against Palestinians lead to escalation rather than incapacitation, and that Palestinians are in fact reacting to Israeli behavior.[6]

We expand on that literature by describing full response curves which capture equilibrium strategies and solve for equilibrium levels of violence. In a technical extension we find, relative to that literature, that a VAR approach suffers from bias due to temporal disaggregation, which can become extreme if response lags are much longer or shorter than the frequency of

[6] Golan and Rosenblatt (2011) comment on Haushofer, Biletzki, and Kanwisher (2010) showing that the empirical response function of Jaeger and Paserman (2008) can be misleading, and demonstrate the sensitivity of their results to choosing lag lengths jointly or individually. Finally, Asali, Abu-Qarn, and Beenstock (2017) revisit Jaeger and Paserman (2008) focusing on modeling the problem nonlinearly. They show the sensitivity of the Jaeger and Paserman (2008) results to modeling choices, and that VAR innovations are not normally distributed, leading them to estimate nonlinear vector autoregressions. They conclude that both sides were reacting to the other's violence, as did Haushofer et al (2010).

measurement (in “calendar” time). We illustrate through simulation that VAR may overestimate (or underestimate) duration of response, potentially allowing false inference about impulses causing escalation (Klinenberg et al., 2024).

This study also builds on a larger literature addressing the logic of terrorism and counterterrorism in Israel-Palestine (Kydd and Walter 2002; Evron 2005; Berman and Laitin 2008; Jaeger and Paserman 2009; Rubin 2011; Dugan and Chenoweth 2012; Jaeger et al. 2012 and 2015; Benmelech et al. 2015; Sobelman 2017; Durante and Zhuravskaya 2018; Nanes 2019). Another large literature documents the effects of violence and of Gaza’s economic isolation (Getmansky and Zeitzoff 2014, Etkes and Zimring 2015; OCHA, 2016). We complement a smaller, relevant literature on proxies and agents (e.g. Padró i Miquel and Yared, 2012, Berman et al., 2019) by modeling the sides as unitary actors.

We proceed from deterministic to stochastic response curves in Section Two, explain our data in Section Three, and how we code it in Section Four. Our approach to constructing a predicted damage measure using casualty data is covered in Section Five, a prelude to estimating response curves in Section Six. Section Seven explores the effects of major operations in Gaza, as well as that of the Iron Dome missile defense system. Section Eight concludes.

2. Response Curves and Deterrence

We begin by analyzing escalation and de-escalation in a two-sided repeated game within a deterministic model. We then generalize to the stochastic case in which opponents adopt mixed strategies, which the data will insist on. (Our aim throughout is the simplest possible deterrence model consistent with our data.) Putting aside for a moment the welfare calculations that underly response curves, we start by defining terms and characterizing equilibria.

Two sides A and B suffer damage inflicted by the other, D^A and D^B, respectively. A’s response curve is illustrated in Figure 1: A responds to damage suffered D^A (measured on the horizontal axis) with $D^B = R^A(D^A)$ (on the vertical).

Figure 1: Response Curve

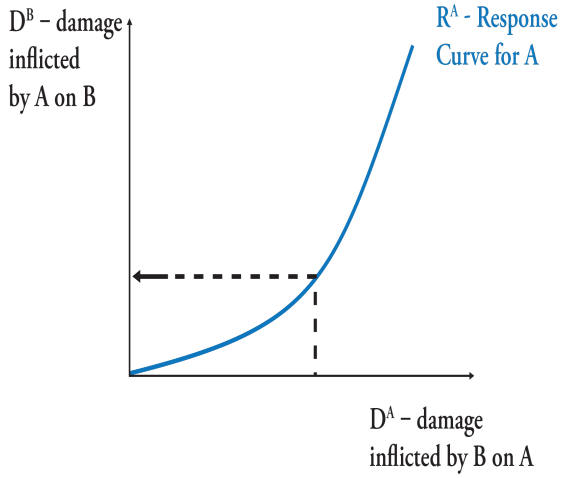

B responds with $D^A = R^B(D^B)$, as illustrated by the additional green line in Figure 2 (bearing in mind that B suffers damage measured on the vertical axis and responds with damage measured on the horizontal axis). In the technical terms of modern game theory, they are engaged in a sequential, noncooperative, repeated, negative-sum, full-information game, with full commitment.[7]

Figure 2: De-Escalation to Equilibrium

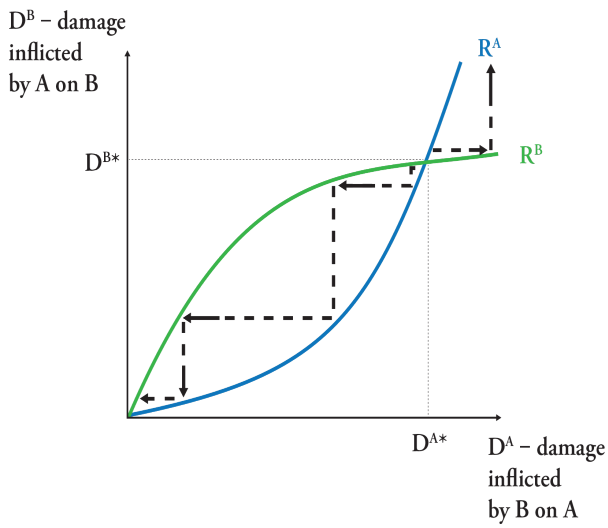

In Figure 2, both A and B have deterrent postures, by which we mean that their response increases in damage incurred, i.e., imposing increasing costs in response to increased damage.

[7] Ruling out simultaneous play and incomplete information simplifies our analysis. For precise implications of the full set of assumptions see Dixit and Skeath (1999, Chapter 2).

We assume that each side derives response curves strategically to maximize a hypothesized welfare function subject to technological and political constraints. While we observe damage, we don't directly observe the underlying disutility associated with it in each side's welfare function. Nor do we observe the costs each side incurs in inflicting damage on the other, or their other constraints, treaty obligations, or other considerations involving third parties (allies or enemies) – all of which could influence the optimal response curve. Analytically, this is analogous to observing supply and demand curves without observing production or utility functions, with the added consideration that a response curve is a strategic choice that takes into account the opponent's response curve. We will make assumptions below on the welfare function as necessary to generate predictions.[8]

2A. Escalation and de-escalation

A sequence of responses and counter-responses is de-escalating if $R^B(R^A(D^A)) < D^A$ and $R^A(R^B(D^B)) < D^B$, at some point (D^A , D^B), i.e., a sequence of attacks that decline in damage. Those two conditions are equivalent (assuming differentiable response functions) since both require that the product of derivatives is less than one [$R^{B\prime}(\cdot) \times R^{A\prime}(\cdot) < 1$]. A sequence of attacks escalates if the opposite occurs, i.e., $R^{B\prime}(\cdot) \times R^{A\prime}(\cdot) > 1$.

The serrated lines in the southwest corner of Figure 2 illustrate a *de-escalating* sequence, which seems roughly equivalent to what the IDF and Hamas describe as an episode on their social media. Where does the de-escalating sequence end? At a fixed point (D^{Af} , D^{Bf}) with $R^B(R^A(D^{Af})) = D^{Af}$ and $R^A(R^B(D^{Bf})) = D^{Bf}$, we have an equilibrium, such as the point (0,0) in Figure 2, where the two response curves intersect. That equilibrium is stable if the sequence de-escalates for D^A just above D^{Af} (and escalates just below). In Figure 2, (0,0) is a *stable* fixed point because the slopes of the two response curves imply de-escalation for levels of damage just above zero. (More precisely, the green line is steeper than the blue, since $R^{B\prime}(\cdot) \times R^{A\prime}(\cdot) < 1$, so $R^{A\prime}(\cdot) < 1 / R^{B\prime}(\cdot)$).

[8] We do assume that each side is a unitary actor, an assumption relaxed in Nanes (2019) discussing the same conflict, and in Dixit and Skeath's (1999) treatment of the Cuban missile crisis, pp. 446-450. More specifically, we assume that even if Hamas is acting as an agent incentivized by an Israeli principal to suppress attacks by other militants, as in Nanes (2019), the welfare function of that non-unitary system in Gaza would still generate an upward-sloping response curve. We thank Matthew Nanes, our discussant, for helping us clarify this point.

In contrast (D^{A*}, D^{B*}), the crossing point of the two response curves in the northeast (where the 'tail of the fish' begins), is an *unstable* equilibrium, since the slope of the blue line exceeds that of the green. For example, an attack on A with damage exceeding D^{A*} would invite a responding attack on B that exceeds D^{B*}, and an *escalating* sequence of attacks would follow, as illustrated by the serrated lines northeast of the intersection point.

Stable response curves in repeated episodes (i.e., sequences of attacks and responding counterattacks) would imply a dynamic equilibrium of repeated, de-escalating episodes in positive intervals bounded by the origin and (D^{A*}, D^{B*}). Damage incurred at or exceeding D^{A*} or D^{B*}, respectively, would imply escalation. In this sense escalation is an implication of the shape of *both* response curves; in Figure 2 an attack escalates the system by inducing a response northeast of the unstable equilibrium (D^{A*}, D^{B*}). In summary, the positions of equilibria, their stability, and escalation or de-escalation are a function of both curves' position and shape. Figure 2 also illustrates how the convexity of the two response curves implies the existence of the unstable equilibrium in the northeast, motivating our interest in concavity and convexity of response curves (beyond our interest in intercepts and slopes).

Why would response curves have these shapes rather than others? Maximizing welfare (subject to constraints) might result in *convex* response curves which would better deter high levels of damage, by raising the marginal cost of inflicting increased damage for an opponent. One reason is a lack of capacity. Another is that disproportionate violence may be unacceptable to domestic or international audiences, making a convex response curve not credible. A third reason is the riskiness of accidentally landing to the right of the northeast equilibrium in Figure 2, doomed to escalation. Finally, the figure also illustrates that increasing convexity comes at the cost of making accidental escalation more likely. Imagine that A shifts to an even more convex response curve (though still anchored at the origin), by everywhere increasing the slope of the blue line. Doing so would shift the new northeast equilibrium to the southwest, narrowing the de-escalation interval, and *increasing the chance of an attack accidentally landing in the unbounded escalation interval* east of that equilibrium. Of course, that thought experiment assumes no strategic response by B, who might be scared into reducing convexity (or slope) –to prevent escalation, or might react symmetrically *–further contracting the de-escalation space* (again assuming that B is not constrained by capacity, domestic or international pressure).

Key to our approach is the tension between short- and long-run damage suffered. Either side could immediately de-escalate to the low violence equilibrium, (at (0,0)), gaining relief from violence this round. Yet that would imply a flat response curve—which in the next round will fail to disincentivize attacks by the opponent (by not imposing costs). This inherent tension is known as the commitment problem in the literature on repeated games (e.g., Dixit and Skeath, pp. 292-306). The core implication is that sides seeking to maintain credibility will keep response curves stable between rounds. We will test for that stability empirically.

Relevant to escalation management are two further empirical questions: First, what are the upper bounds of the de-escalating interval (D^{A*}, D^{B*}), beyond which both sides are doomed to escalation?

Second, is the stable, low-violence equilibrium at (0,0)? In the literature on deterrence between two nuclear powers, the origin was an aspirational equilibrium of great importance, not least because de-escalating response curves at low levels of damage (presumably using conventional weapons) could not be assumed. Any deviation from the origin, which thankfully never occurred during the Cold War, might have been catastrophic.

In contrast, in the current discussion of integrated deterrence involving conventional weapons, we frequently see attacks and damage: for example, cyberattacks, election interference, sanctions, and other damaging violations of international law and sovereignty. Attacks might occur in equilibrium depending on the shape of the response curves. A deterrent posture implies only that response curves are increasing in damage suffered without any other restriction on their shapes. This may lead to one stable equilibrium at the origin, but it could also lead to a stable equilibrium with positive violence (as we will estimate below), and even to multiple stable equilibria.

Figure 3 illustrates response curves that generate (a stable) equilibrium with incomplete deterrence, by which we mean that at least one side suffers (nonzero) damage in equilibrium. In this example, response curves are concave, and do not go through the origin. The vertical intercept of A's response curve is above the origin, and the horizontal intercept of B's response curve is to the right of the origin. This implies that A and B carry out "unprovoked" attacks, as they respond with positive damage even when they suffer zero damage. In this scenario, incomplete deterrence occurs: the resulting equilibrium involves damage to both sides. A secondary implication is that there is only one equilibrium.

Figure 3: Equilibrium with Incomplete Deterrence

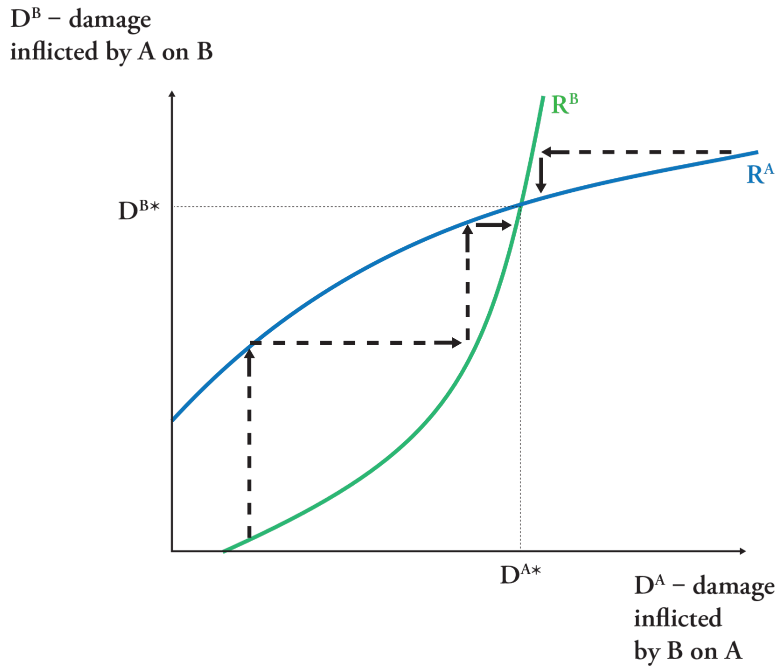

We should emphasize that there is no completely peaceful equilibrium in this case. If the two sides find themselves at the origin, one or the other would initiate an unprovoked attack, leading to an escalating sequence that ends at the same violent equilibrium, as illustrated by the serrated lines southwest of the equilibrium. The result is stable (since the green line cuts the blue from below in the neighborhood of the equilibrium), though violent.

An equilibrium with incomplete deterrence invites the question of policy options. We turn next to actions the sides might take to shift such an equilibrium by influencing an opponent's response curve.

2B. Deterrence by Denial and by Resilience

The deterrence concept discussed so far operates by imposing costs –tracing an upward-sloping response curve. It is distinct from deterrence by denial (the basic logic of conventional warfare), in which one attempts to reduce or degrade the opponent's capacity to attack (by destroying munitions, assassinating leaders with irreplaceable skills, or eroding public support for continued attacks, for example).

Figure 4 illustrates. Imagine that A destroys some capacity of B. The effect on B's response curve depends on B's underlying welfare and cost functions, which we are now forced to consider. A fairly general assumption is that destroying B's capacity will raise B's marginal cost of damaging A (more than it might increase B's marginal return from damaging A) at all levels of D^B. For instance, if Israel destroys some Gazan rockets with an airstrike, it may not fully preclude the possibility of a future high-damage attack, but could instead force Gazan militants to use more mortars (or put themselves at risk by moving mortars closer to the border fence) to achieve the same levels of damage. If so, A will have succeeded in shifting B's response curve leftwards to lower levels of damage, as shown by the serrated green line. The resulting equilibrium (to the bottom left, where the serrated response curve R^{B}' meets R^A) will occur at lower levels of damage for both A and B, when compared to the previous equilibrium (though the welfare effects are presumably higher for A and surely lower for B).

Figure 4: Deterrence by Denial, and by Resilience

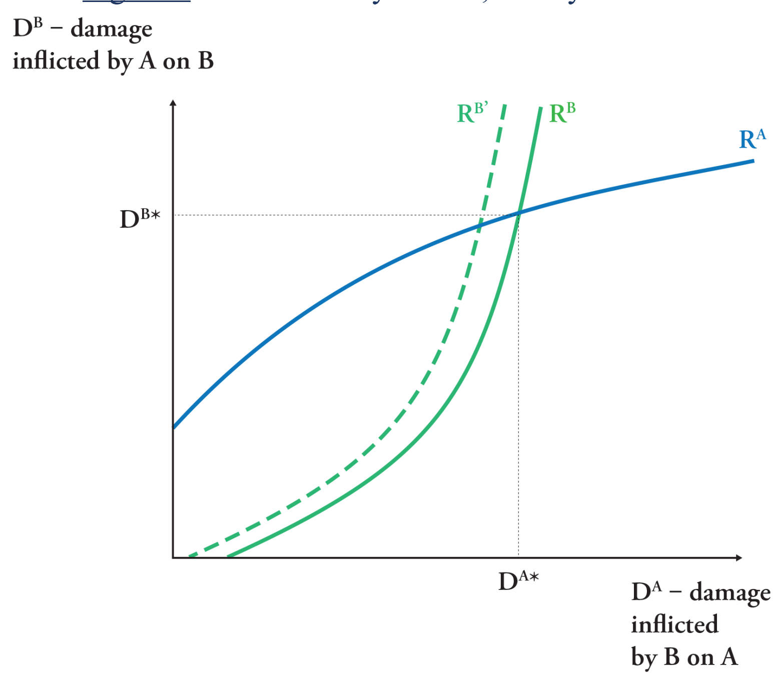

Our data include examples of three major Israeli military operations and other large actions that destroy Gazan capacity, so the hypothesis that they shift the Gazan response curve to a lower level of violence is testable. Those operations are often criticized as merely "mowing the lawn," in the sense that destroyed capacity might be restored (just as grass grows), yet in our model, the operation might be nevertheless pursued by a rational actor if a long post-operation period of less violent equilibrium, even if limited, justifies the one-time cost of the operation (see also Inbar and Shamir 2014; Gibilisco 2023).

Deterrence by resilience, by contrast, reduces or even prevents damage associated with attacks initiated by an opponent. This should be the case for the "Iron Dome" air defense system, which arguably makes Israel more resilient and may affect Gazans' damage function. Under the fairly weak assumption (about the underlying welfare function) that the increased resilience of A reduces the marginal benefit to B (through cost imposition) of attacks perpetrated at all levels of damage, we can predict how B's response curve will shift. In the diagram, if A increases resilience, that would shift B's response curve to lower levels of damage (to the left), as did deterrence by denial (above). The resulting equilibrium is again illustrated in Figure 4 at lower levels of net damage. As the diagram illustrates, increasing resilience has very similar effects on equilibrium outcomes as does destroying capacity –though it is achieved through defensive rather than offensive operations.

An additional prediction of a general strategic model is that a shift in an opponent's response curve may generate an adjustment of one's own. Under fairly general assumptions, both deterrence by denial and by resilience (shifts of an opponent's curve) will complement deterrence by cost-imposition (slope of one's own curve), so that a shift to the left of an opponent's response curve (i.e., to a less aggressive posture) will induce a shift to the left of one's own (more aggressive); the logic being that the benefits of cost-imposition remain the same, while the costs (in expected retaliation) have declined. A follow-on implication of that additional prediction is that major military operations (denial) and Iron Dome (resilience) unambiguously shift the subsequent equilibrium to lower damage for Israel. Whether the new equilibrium has lower damage for Gaza (as illustrated in the diagram) is theoretically ambiguous (because the blue response curve might shift as well).

Estimation will allow us to test the general applicability of this model in two important ways. To get an interval with de-escalation, the shape of response curves must yield a fixed point (i.e., a stable steady state) at the low end of the interval. Additionally, for these episodes to constitute a dynamic equilibrium as a group, the response curves must be *stable across episodes as well as within*. Note that the stability of response curves constitutes refuting evidence for a competing model in which attacks reduce opponent capacity (deter by denial).

2C. Stochastic Response Curves

The available data show that responses are stochastic. The Israel Defense Forces (IDF) do not respond to 73 percent of initial attacks from Gaza during our sample period, while Gazan forces do not respond to 77 percent of initial Israeli attacks. Moreover, the intensity of response shows considerable variance, even conditional on damage suffered. That behavior motivates a model with mixed strategies—i.e., responding in probability, and with stochastic damage inflicted.

Why mixed strategies? Intuitively, mixed strategies would be preferable to always responding if there are fixed costs to responding at all, as a means of maintaining deterrence against relatively low levels of damage incurred (if marginal costs of response exceed marginal benefits). Mixed strategies also solve the dilemma of achieving deterrence by creating an expectation of costs imposed without being predictable about when (or how or where) a response will come, forcing the opponent to incur excess mitigation costs (e.g., taking shelter).

In the stochastic case, both A and B play mixed strategies. They respond with probabilities $P^A(D^A)$ and $P^B(D^B)$, respectively. That is, both sides respond to damage (suffered) with an expected amount of damage inflicted. That approach is coherent if we generalize welfare maximization to expected welfare maximization, assuming risk-averse opponents.

If the probability of response increases in damage suffered, and response curves (conditional on some response) are increasing in damage suffered, then expected response curves, such as $E(D^B) = P^A(D^A) \times R^A(D^A)$, will also increase in damage suffered. As in the deterministic case, the resulting expected response curves should arguably adopt a deterrence posture, but their exact shape is determined by the maximization of welfare subject to (political and capacity) constraints. The tension between short- and long-run returns to retaliation (described above) carries over to the stochastic case, as do the discussions of stability, de-escalation, and unique equilibria, though now in expected damage incurred rather than in deterministic damage.

The stochastic analysis is somewhat deceiving in its simplicity. An equilibrium in mixed strategies appears as a point in expected damage space on a graph, but actually represents a bivariate distribution of damage suffered by A and B concentrated around that point –which reflects only the *expected values* of damage.[9] Anticipating our empirical results, we will estimate a stable, stochastic equilibrium—the first we know of in this literature on conflict.

---

[9] Technically, those distributions are truncated at the horizontal and vertical axes, which is how we will implement them in estimation.

3. Data

Our study of mutual deterrence is enabled by a unique dataset describing in remarkable detail the repeated violent interactions between the IDF and Gazan militants. These data are unusual in their recording frequency, the richness of incident reports, and the availability of other supporting data, including from the social media of the perpetrators, to validate and assist in interpretation and coding.

We build our main dataset using daily security reports of all attacks in Gaza recorded by the United Nations Field Security Office. The reports document attacks at up to five-minute precision, along with the governorate where they occurred.[10] We begin analyzing the UN reports on June 15th, 2007, the day Hamas seized control of the Gaza Strip from Fatah security forces. In our analysis sample, we also omit the period after December 15, 2017, when Gazans began regularly approaching the border fence, violating the buffer ("no go") zone that Israel claimed in order to secure the fence. While those events began as nonviolent, they eventually included acts of violence by young men, endorsed by all the militant groups, which drew Israeli warning shots, exchanges of live fire, and often fatalities on the Gazan side.[11] We exclude them for two reasons: first, in retrospect, those actions contested the buffer zone, which would be used to launch the attacks of October 7, 2023; second, they represented a heightened willingness to take on risk by young men who are classified as civilians (rather than militants) in United Nations reports. Taken together, these events may reflect increased willingness by Hamas and other militant groups to incur casualties, or possibly a change in strategy, presenting a challenge that we defer to future research.

Figure 5 reproduces the security report of November 20, 2010, a particularly violent day we have chosen to convey the richness of the data (text color appears in the original). Reports typically summarize an attack in two or three sentences. On November 19, Gazan militants fired mortars at Southern Israel at 12:30 and 13:50; the Israeli Air Force (IAF) retaliated by firing four missiles at different targets in the Gaza Strip between 15:15 and 15:40; Gazan militants counter-retaliated by firing one mortar at 19:25, and the IAF attacked again at 23:30.

---

[10] Gaza has five governorates. Attacks in two of them, North Gaza and Gaza, are reported together by the United Nations in the earlier years of our sample, so we combine those in our study.

[11] UN Daily Security Update of 15 December 2017 documents four demonstrations at separate locations adjacent to the security fence in three governorates on the same afternoon (of 14 December). Those demonstrations prelude the "March of Return" confrontations on the border fence the following Spring. See also the report on the "March of Return" in the newspaper Ha'aretz, on March 31, 2018 (Khoury et al., 2018).

Figure 5: Security Report for November 20th, 2010

SECURITY UPDATE

20 November 2010

UNDSS GAZA

SECURITY PHASE: IV **SUNSET: 1650 hrs**

AREA SUMMARY (24 hours)

1. **RAFAH:**

- **19 Nov 2330 hrs an IAF F16 fighter fired 2 missiles targeting the smuggling tunnels area, opposite Yebna camp, west of Rafah, no immediate reports of injuries. Israeli sources reported that the Air Force bombed smuggling tunnels in the southern Gaza Strip Friday night and direct hits were identified, the IDF said in a statement. No injuries were reported.**

2. **KHAN YOUNIS:**

- **19 Nov 1230 hrs PRC militants fired 3 mortar shells from east of Al Qarara, east of Khan Younis, towards Eshkol Regional Council in direction of Ein Hashlosha. No reports of injuries or damage. Israeli sources indicated that Palestinians fired three mortar shells towards the southern Negev region. They landed in an open area in the Eshkol Regional Council. No reports of injuries or damage.**
- **19 Nov 1520 hrs an IAF F16 fighter fired 1 missile targeting an open area east of Bani Suhaila, east of Khan Younis. No reports of injuries or damage.**
- **19 Nov 1540 hrs an IAF F16 fighter fired 1 missile targeting an Islamic Jihad training base in the former Neveh Dekalim settlement west of Khan Younis. No injuries reported.**

3. **MIDDLE AREA:**

- **19 Nov 1515 hrs IAF F16 fighter fired 2 missiles targeting a house under construction east of Deir El Balah, 150m west of the dumping site. Severe damage reported to the house and 3 persons were wounded.**
- **19 Nov 1925 hrs unidentified militants attempted to fire 1 HMR from northeast of Nuseirat camp towards the western Negev. However, the rocket exploded prematurely. No faction claimed responsibility. No injuries or damage reported.**
- **20 Nov 0530 hrs unidentified militants fired 4 mortar shells from east of Abu El Ajeen, east of Deir El Balah towards Kissufim military base. No injuries or damage reported.**

4. **NORTH AND GAZA:**

- **19 Nov 1350 hrs unidentified militants fired 5 mortar shells from north of Beit Lahia towards Ashkelon Coast Regional Council in direction of Zikim. However, 1 shell dropped-short and exploded near the security fence north of Beit Lahia. Israeli sources indicated that 4 mortar shells fired from the northern Gaza Strip landed at the Ashkelon Coast Regional Council landed in an open area. No injuries or damage were reported.**

We record the attacker, target, type and quantity of munition, and casualties associated with each attack. For the Israeli side, this entails the number of airstrikes, small arms fire incidents, shelling, and incursions. For the Gaza side, reports include the number of small arms fire incidents, grad rockets, other rockets, and mortars. Deaths and injuries are recorded for Gazan civilians, Hamas operatives, other militants, and Israelis. The reports suggest that small-arms fire is not always used to directly inflict damage. For example, the Israeli Defense Force would fire into the air at border fences to disperse large crowds, and the Israeli Navy would fire across the bow of Gazan fishing boats that drifted into restricted waters. We remove such situations by only coding Israeli small-arms fire accompanied by an incursion or resulting in a casualty. (Our findings are robust to including or excluding all Israeli small-arms fire.)

Figure 6 presents daily tallies of Israeli and Gazan attacks between June 15th, 2007 and December 31st, 2018. In total, Israel launched 10,269 airstrikes, shelled Gaza 22,659 times, performed 1,379 incursions, and fired small arms 1,285 times (implementing our small arms inclusion rule discussed above). Gazan militants fired 10,207 rockets, 6,225 mortars, and 982 grad rockets into Israel. Gazan militants also fired small arms 355 times. The figure illustrates the high frequency of violence during our sample period. Out of 4,218 days included in the sample, Gazan militants performed some form of a violent attack in 33.1% of them (1,396 days) and Israel attacked in 45% of them (1,894 days). This violence caused 38 Israeli deaths (374 Israelis were injured). On the Gazan side we observe 3,522 fatalities, and 17,281 individuals were injured. Out of all Gazan fatalities, 748 of them are civilians and the remainder militants (2,070 injured Gazans are civilians).[12]

These numbers mask a great deal of variation in violence over time, as depicted in Figure 6. The figure shows a relatively low number of daily attacks, together with extremely violent periods occurring right before and during major confrontations between Israel and Hamas. These include the three main Israeli military operations during our sample period: Cast Lead (December 27th, 2008 until January 18th, 2009), Pillar of Defense (November 14th, 2012 until November 21st, 2012), and Protective Edge (June 12th, 2014 until August 26th, 2014). Gazan attacks are more evenly distributed over time, whereas Israeli attacks substantially increase during major operations. Note that violence substantially decreases after those operations vis-à-vis the level of violence before the operations. Within 30 days before those operations, Gazan militants attacked Israel 547 times and Israel attacked Gaza 267 times. The total number of attacks decreases to 258 from Gaza and increases to 383 from Israel for the 30-day windows after those operations.

[12] Among militants, 227 fatalities were Hamas with the remainder belonging to other militant groups, including participants at violent protests at border fence (affiliation not listed, but protests may be organized by Hamas).

Figure 6: Gazan and Israeli Attacks 2007-2018.

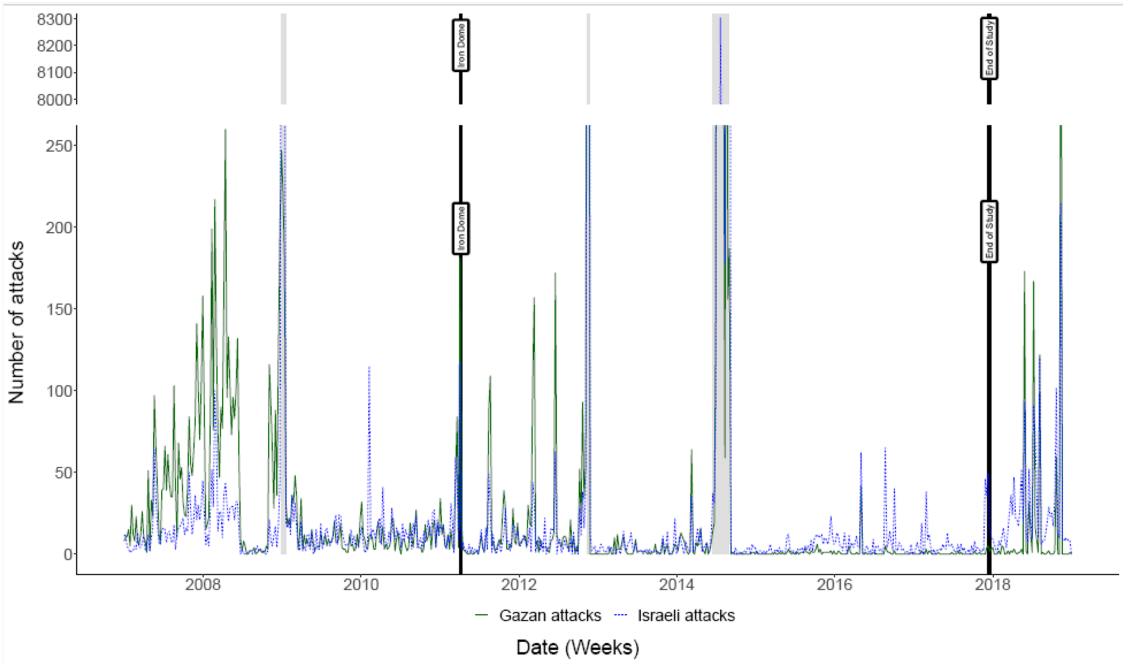

Note: Vertical axis break at 250 Attacks

We omit the periods of the three major Israeli military operations in Gaza from our analysis sample for three reasons: First, descriptions of those operations characterize them not as escalatory responses to a particular attack (or sequence of attacks), but instead as pre-planned campaigns to reduce militant capacity. Indeed, the Israeli government announced that the purpose of operations in Gaza was to hamper Hamas' military capabilities, such as by demolishing tunnels or destroying munitions depots (Inbar and Shamir, 2014). Second, Gazan militants' response would have been constrained by operational ability, being under attack (this is evident in Figure 6). Third, even if those attacks and counterattacks fit on some response curve, their sequence of response and counter-response, and the available recording of that sequence, would have been disrupted by the very intensity of conflict around these operations.

4. Actions and Episodes

Even though the sides describe their attacks as retaliatory, they sometimes inflict a sequence of attacks before suffering a response. We address this by aggregating attacks into *actions*, where an action is a sequence of attacks executed by the same side with no attack executed by the other side between them.[13] Choosing not to respond is also an action –a *zero damage action*.

Actions, in turn, occur in *episodes*, which we will think of as rounds in a repeated game. (We provide evidence below that response curves are stable from episode to episode, reinforcing the idea that both sides treat these as rounds.) One challenge to coding is that the protagonists don't announce when an episode has ended, they simply stop retaliating. Our coding rule is to end an episode if there ensue 48 hours of *lull* (no attacks). We use a 48-hour lull rule to aggregate actions into episodes for two main reasons. First, the IDF and Gazan militants tend to retaliate within one or two days of the initial attack (as reported by their social media), with few exceptions.[14] Second, speedy retaliation makes sense because each side wants others to correctly interpret its action as a retaliation, i.e., as the price of causing damage, establishing this causal linkage (in the minds of internal and external audiences). Moreover, both sides have the technological capability to retaliate within hours of an attack. Israel keeps fighter jets ready to scramble, drones hovering in the skies over Gaza, and tanks around its borders. According to our data, these systems oftentimes spot and strike militant rocket crews even as they set up to launch. Similarly, as Haushofer et al. (2010) established, the technology for launching mortars and rockets allows Gazan militants to retaliate within hours of provocation.[15] An *episode* is thus defined as a group of actions bookended by lulls. An episode ends when one side does not retaliate, which we code as a zero-damage action. We code lulls as two zero damage actions for every 48 hours elapsed between (nonzero) actions.

The UN records in which of four governorates attacks occurred, as we saw in Figure 5. We present three approaches to aggregate actions and episodes across governorates, each relying on different assumptions. The first is to assume that all actions in the Israeli-Gazan conflict are

[13] Marlin-Bennet et al (1991) refer to these as 'moves' and 'turns,' in an analogous coding.
[14] Abrahams et al (2019) report data on retaliation from the social media accounts of both the IDF and Gazan militants between April 2009 and July 2016, collected by the Meir Amit Intelligence and Terrorism Information Center (http://www.terrorism-info.org.il/). Of 305 retaliatory IDF airstrikes in that sample they find that at least 95% occurred within 2 days of initial attack. Of 26 retaliatory militant rocket attacks, 22 (85%) occurred within two days.
[15] This contrasts with suicide attacks launched from the West Bank during the Second Intifada, which potentially required weeks of planning (Jaeger and Paserman, 2008).

responses to the other side's actions, regardless of location (effectively assuming that Gaza is one conflict zone). We refer to this as the *broad* specification. A concern with this approach is measurement error in estimating response curves' slopes because we are likely estimating some actions as responses to damage suffered elsewhere, when they were not. For example, this aggregation may mistake an incursion in Rafah governorate as a response to mortar fire coming from Khan Yunis, when those are independent attacks. Our second approach is to assume that actions are only responses to attacks within the same governorate. We refer to this as the *local* specification. This aggregation method may misrepresent actions responding across governorates as being unprovoked attacks beginning an episode. Econometrically, it might particularly bias the response curves' intercept because we are failing to attribute responses to actions outside the governorate. The final approach is a mix of local and broad specifications, which we refer to as *hybrid.* This combines episodes from the *local* specification together if an actor uses the same non-small arms munitions in two governorates within one hour of each other –our logic being that small arms fire may be controlled at the local level, whereas rules of engagement for more lethal munitions require approval at a higher level. We prefer this *hybrid* specification because we think it most accurately captures cross-governorate responses, so we use it throughout the remainder of the paper.[16] Results below are robust to using broad or local specifications instead.

Table 1 reports on actions (omitting periods of major operations) for those three approaches to coding episodes. Focusing on the hybrid approach, Israel's attacks aggregate to 7,978 actions, and Gazans' attacks to 7,977. Most actions are zero damage. Israel chooses to carry out an attack in 2,124 of its actions (27%). Similarly, Gazan militants carry out attacks in 23% of their actions (1,827/7,977). This motivates our interest in mixed strategies. Most Israeli actions consist of shelling, airstrikes and incursions, which cause on average 0.11 Gazan fatalities and 0.41 injuries, per action. Most Gazan actions consist of mortar and rocket attacks. They average very few Israeli fatalities (0.003) per action, but result in 0.03 injuries. The broad coding (compared to the hybrid) records fewer actions, because it aggregates more attacks across governorates into a single action, while the local coding disaggregates those same attacks into more actions. For that reason, broad episodes also record higher average fatalities and injuries, and local episodes less.

Recall that we aggregate actions bookended by lulls into *episodes*, which we will treat as rounds of a repeated game. Table 2 reports episodes' characteristics. The main message of the

[16] See online appendix for a detailed explanation of the hybrid aggregation method.

table is how much episodes vary. Focusing on the hybrid episodic coding, duration ranges from five minutes (zero days) to 121 days, and from a single action up to almost 500 (248 Gazan + 247 Israeli). That said, most episodes are short-lived, with over half lasting less than five hours. The median lull length is 6.6 days. Hence, a typical week during our decade-long sample period has one violent episode that includes one Israeli action and/or one Gazan action.

Just less than half of the episodes consist of attacks by only one side. Fourteen percent of episodes have only Gazan attacks and 32% have only Israeli attacks. Most of these single-action episodes consist of one attack (lasting for our minimum coded duration of five minutes), and do not cause any damage to the other side. Single-action episodes with Gazan attacks consist mostly of rocket or mortar attacks, often landing on open fields, without Israeli retaliation. Single-action episodes composed of only Israeli attacks tend to be small arms fire during incursions.

Treating the conflict as a sequence of episodes and lulls allows us to estimate response curves, and test their key properties: specifically slope, concavity, and stability.

5. Measuring Damage

Estimating response curves requires a measure of damage per attack that somehow combines different characteristics of harm, and associates them with different types of attacks. The assumption that we can compare harm caused by different weapons is key to our integrated deterrence approach. Characteristics of harm range from the horrible to the more mundane, which in principle include not only fatalities but also physical injuries, damage to property, psychological trauma, as well as avoidance and mitigation costs. Fully measuring damage is impractical, so we approximate using casualties (fatalities and injuries) reported after each attack, experimenting with different weights on fatalities and injuries.

We create an aggregate measure of predicted damage using as weights coefficients from linear regression of casualties on types of munitions used by an opponent. Specifically, we estimate the following linear models to generate weights, where *i* denotes an attack.[17]

Damage to Gaza (D^G):

$$D_i^G = \beta_1 small\ arms_i + \beta_2 incursions_i + \beta_3 small\ arms\ fire_i * incursions_i + \beta_4 shellings_i + \beta_5 airstrikes_i + \epsilon_i \quad (1)$$

[17] Recall that attacks are documented at five minutes intervals. A single attack often consists of a barrage using different types of munitions simultaneously.

Damage to Israel (D^I):

$$D_i^I = \alpha_1 small\ arms\ fire_i + \alpha_2 mortars_i + \alpha_3 rockets_i + \alpha_4 grads_i + \epsilon_i \quad (2)$$

We allow no intercept, assuming that casualties cannot occur without an attack.[18]

Table 3 reports estimated coefficients for equations (1) and (2), exploring four alternative ways to measure damage using casualties. These are estimated using all attacks causing casualties (see the discussion of Table 1). Column 1 uses exclusively fatalities as our measure of damage; Column 2 assigns weight one to fatalities and 0.2 to injuries; Column 3 assigns weight one to fatalities and .5 to injuries; and Column 4 uses equal weights.

Estimated coefficients are quantitively proportional for different weighting schemes. All the coefficients are positive, as expected. Regardless of the weights used for casualties suffered by Gaza, small-arms fire has a higher coefficient than do incursions, airstrikes, or shellings, with the relative size of coefficients similar regardless of the weight placed on injuries. Israeli small-arms fire during incursions causes far less fatalities and injuries, so we estimate a separate coefficient for it, to achieve a more accurate weight. In terms of damage suffered by Israel, small-arms fire again has the highest coefficient, followed by rockets and then mortars, regardless of the weight placed on injuries. The coefficient on Grad missiles is not statistically significant, perhaps because a large number of these missiles are intercepted by the Iron Dome system. (Though we do not find missiles to be any less lethal following the introduction of Iron Dome –in unreported results.) We return to discuss the Iron Dome missile defense below. Coefficients on small arms fire, rockets, airstrikes, and incursions are quite precisely estimated.

Given that the results (below) are qualitatively robust across weighting schemes, we use the fitted values from the second column as the preferred aggregate damage metric for each attack, as it counts injuries as well as fatalities, but with less weight.[19]

---

[18] All reported casualties are associated with at least one specific munition, so a statistically significant intercept would likely reflect a specification error in estimation.

[19] Recalling the discussion above of events contesting the "no go" zone in December 2017, one reason to drop data beginning in that period is the possibility that the weight Hamas assigned to civilian casualties had shifted, which would confound our measurement.

6. Estimating Response Curves

6A. Functional form and hypotheses

We choose simple functional forms to estimate response curves, to preserve precision in estimation, and allow for possible serial correlation. The Israeli response curve (D^G), which measures damage to Gaza, is given by

$$D^G_a = \beta_0 + \beta_1 D^I_{a-1} + \beta_2 (D^I_{a-1})^2 + \varepsilon^I_a \quad (3)$$

while the Gazan response curve (D^I), which measures damage to Israel, is given by

$$D^I_a = \beta_3 + \beta_4 D^G_{a-1} + \beta_5 (D^G_{a-1})^2 + \varepsilon^G_a. \quad (4)$$

where the subscript 'a' denotes an action. Response curves are quadratic in damage suffered to allow for concavity or convexity, recalling the importance of second derivatives (Section 2).[20]

The strategic interaction between sides may give rise to an unusual cross-correlation of errors, as depicted in equations (5) and (6):

$$\varepsilon^I_a = \rho^{IG} \varepsilon^G_{a-1} + \rho^I \varepsilon^I_{a-2} + u^I_a, \text{ for Israel's response curve} \quad (5)$$

$$\varepsilon^G_a = \rho^{GI} \varepsilon^I_{a-1} + \rho^G \varepsilon^G_{a-2} + u^G_a, \text{ for Gaza's response curve,} \quad (6)$$

for actions a, where u^I and u^G are independent errors.

In equations (5) and (6) ρ^{IG} and ρ^{GI} represent the possibility that players respond differently to damage suffered due to actions on the curve (which they might expect) than they do to residual damage, which they might not. Those cross-correlations might be negative, perhaps because they perceive the residual as a mistake and are willing to ignore it, or positive (if deviations require extra deterrence). Or they might reflect specification error on our part. The error terms might also display standard serial correlation, due to omitted variables that persist from action to action, such as weather, availability of munitions, or political climate.

We test five related hypotheses:

1. Expected response curves $E(D^G | D^I)$ and $E(D^I | D^G)$ have a positive slope, consistent with a deterrent posture (as in Figure 2).

[20] Bearing in mind that response is often a zero action (Table 1), we could alternatively estimate separately a probability of response and an intensity of response, i.e., $E(D^G) = P^I(D^I) R^I(D^I | D^G > 0)$, and $E(D^I) = P^G(D^G) R^G(D^G) | D^I > 0)$ (omitting subscript 'a' for readability). A practical problem with that approach is that it is inherently nonlinear. In contrast, our approach seeks to preserve a linear specification in order to deal with possibly serially correlated errors in our time series, in equations (5) and (6). Having chosen a linear approximation, we are mindful going forward of how often (3) and (4) predict negative expected damage.

2. The interaction of response curves yields a stable fixed point at nonnegative values of damage; in other words, a *stable equilibrium.*
   We test sufficient conditions of a stable equilibrium (Section 2): parameters of expected response curves yield de-escalation to the right of the equilibrium and escalation to the left. I.e., a fixed point, $R^G(R^I(D^I)) = D^I$ and $R^I(R^G(D^G)) = D^G$, with de-escalation $R^G(R^I(D^I)) < D^I$ and $R^I(R^G(D^G)) < D^G$, to the right, and escalation to the left (Figure 2).

   2A. If a stable equilibrium exists, is that equilibrium at (0,0), which is to say, complete deterrence in equilibrium? The alternative is *incomplete deterrence in equilibrium*, at positive damage.

3. Response curves are stable (i.e., Markovian) across episodes. We test this by checking if lagged damage and lull length from past episodes (i.e., episodes e-j, for j>0) predict current damage.

4. Response curves are stable (i.e., Markovian) within episodes. We test this by checking if lagged values of damage suffered preceding the last action by the opponent (a-1) predict current damage imposed (i.e., responses to action a-j for j>1).

5. The introduction of the Iron Dome defensive system and major operations do not shift the response curves of Gazan militants. Alternatively, major operations enhance deterrence, shifting Gazan militants' response curve to less violence (presumably by raising the cost of imposing damage on Israel –deterrence by denial), and the Iron Dome achieves deterrence.

6B. Data

In estimating response curves the unit of observation is an action (which may include multiple attacks, as described above). Table 4 reports descriptive statistics for 7,978 Israeli actions and 7,977 actions by Gazan militants recorded during our sample period.

Beginning with Israeli responses, Israel usually takes no action within episodes: its proportion of actions with attacks is 0.27. When Israel attacks, the average damage inflicted is 0.47, recalling that damage is measured as predicted casualties associated with attack type, using as weights the coefficients estimated in Table 3, column 2 (with damage is the equivalent of one fatality or five injuries). We choose predicted damage rather than actual damage suffered, following the logic of a response to *intent* rather than a response to *outcome*, and to capture disruption and mitigation costs in addition to actual casualties (for example, anxiety and moving to shelter), which would occur even in the absence of casualties.

As right-hand side variables in equation (3) we include damage to Israel in the previous action (including zero damage actions) and its square. Note that the median damage suffered is zero because a typical action is unprovoked –as will be true for the Gazan response as well. To check the basic modeling assumption of the repeated game, that episodes are self-contained, we also include the length of the lull preceding the previous episode (measured in hours), and damage to Israel in the previous episode. Those could be motivated by the logic of a delayed response. Note that minimum damage in the previous episode is zero, which occurs when an unprovoked action is not responded to.

Symmetrically, the same set of variables appears in the bottom panel for estimating equation (4), the Gazan response. Gazan militants record 7,977 actions. Gazan militants carry out a positive attack on Israel in 23% of their actions, with average damage inflicted of 0.04 (fatality-equivalents, as above). Damage to Gaza is much greater than that inflicted on Israel, reflecting the asymmetric capabilities of the two sides.

6C. Response Curve Estimates

Figure 7 illustrates our approach to estimating expected response curves. Dots indicate actions as (x, y) pairs of damage: predicted casualties calculated using coefficients from Table 3, column 2. The two panels illustrate the implied expected response functions for Gaza and Israel, respectively, as described in Section 6A above. The plotted lines are from our preferred specification, which we discuss below, with confidence intervals shaded in. Both are upward-sloping in damage suffered, *suggesting a deterrent posture*.

Visually, we don't discern in that cloud of points any thresholds in the response behavior of either side, despite the discussion of "red lines" in commentary on deterrence. That observation reassures us about choosing continuous functional forms for response curves.

Figure 7: Estimating Response Curves

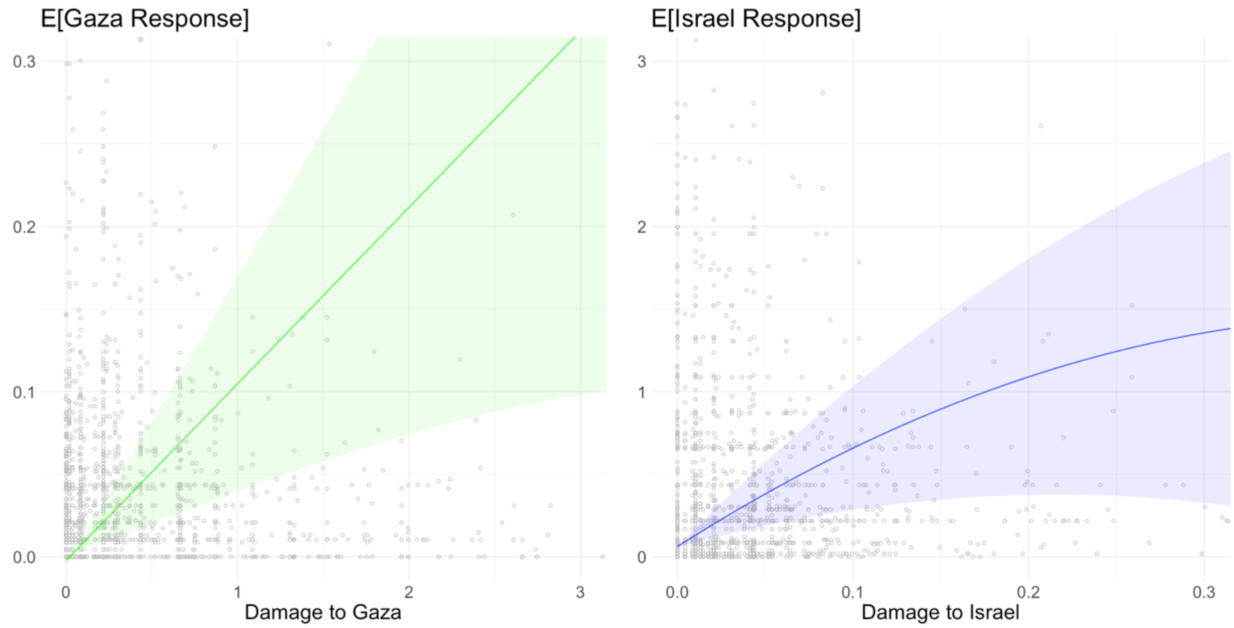

Table 5 describes our response curve estimates (equations (3) and (4) above). They appear in Panel A of the table, labeled Simple Model. Beginning with the Gazan, in column (1), we see that the curve is upward sloping and concave in damage suffered (linear coefficient of .032 and quadratic coefficient equal to -0.0056), evoking the green line in Figure 7. The intercept is positive (.0070), implying a propensity to inflict damage when none is suffered. (Those actions are reflected in the positively valued points on the y-axes in each panel of Figure 7.) The Israeli response curve in column (2) shares those characteristics: positively sloped, concave (linear term of 5.82 and quadratic of -13.39) and positive intercept, like the blue line in Figure 7. (Those two columns are jointly estimated, clustering standard errors by episode.) Note that both y-axes are well populated with actions: positive damage attacks which follow zero damage suffered. This reflects the propensity of both sides to attack in an unprovoked way, or at least to some acts not recorded in UN reports. For Israeli forces, that might be an individual approaching the security fence. For Gazan militants, we can suppose that it might be in response to some other grievance related to the ongoing conflict. All 1,170 episodes begin, in our coding method, with an unprovoked attack. Note also an important contrast between Gazan and Israeli coefficients in their scale, which mirrors the difference we've seen in the data description so far: casualty-based measures of damage for Gaza are an order of magnitude larger. Referring back to the five hypotheses of Section 6A, we have a clear result for the first: both Gazan militants and Israel have an upward-sloping (deterrent) posture.

Before turning to testing stability, and econometric issues, those initial estimates have clear implications for equilibrium –our second hypothesis above. (These basic results will be common to more complex specifications --except as noted below.)

6D. Equilibrium

Figure 8 describes the equilibrium implied by estimated expected response curves of the Israeli side and Gazan militants, taken together. The Israeli expected response curve, in blue, has a positive intercept and a slightly concave positive slope. The Gazan curve, in green (with damage suffered on the vertical axis and damage inflicted on the horizontal) has a positive slope and is very slightly concave.

Since both curves are concave and have positive intercepts, they intersect only once. Because the Gazan line cuts the Israeli from below, that single intersection is a stable equilibrium, de-escalating from the right (i.e., from higher levels of violence) and escalating from the left, as in Figure 3 (of Section 2 above). That is, it satisfies the properties of a fixed point for Hypothesis 2 above. Therefore, the mutual deterrence equilibrium implied by our estimated model results in *a single, stable equilibrium*, ruling out the possibility of escalation illustrated in Figure 2. In other words, almost the entire positive quadrant is a de-escalation zone –above and to the right of the equilibrium.

Figure 8: Equilibrium with Incomplete Deterrence (Simple Model)

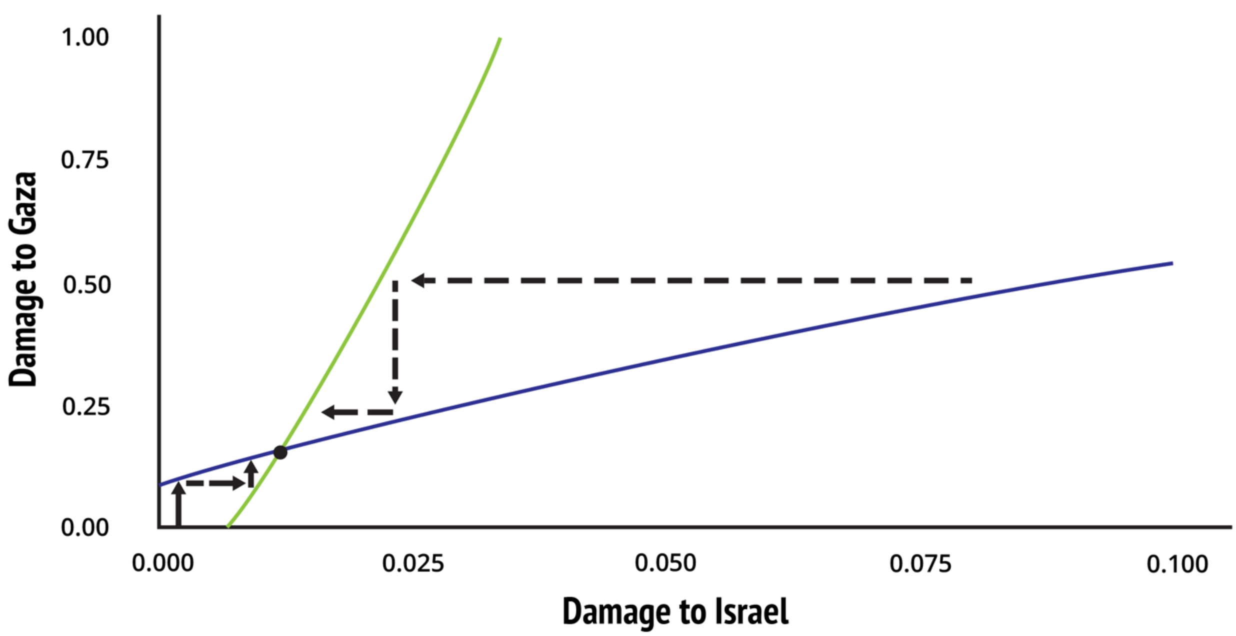

In the remainder of the positive quadrant -below or to the left of the equilibrium, one or the other side increases damage until the exchange converges to the equilibrium, as in Figure 3. In particular, the origin is an unstable point. Zero violence cannot be sustained to the southwest of the equilibrium. E.g., zero violence incurred in Israel induces a response at the intercept of the blue line, which in turn induces a Gazan response, at least in expectation. In this sense, this is an equilibrium of *incomplete deterrence*, as illustrated in Figure 3. The estimated equilibrium is (0.0152, 0.163) in damage units (predicted fatalities and injuries based on analysis in Table 3).

Yet the diagram is deceivingly simple: equilibrium is not a point but a joint distribution reflecting the stochastic response of both sides (centered around the point of intersection). Attacks and actions will generate sequences of damage (within our episodes) that return to this equilibrium in probability, but not deterministically.

To our knowledge, this is the first empirical characterization of a stable deterrence relationship using estimated response curves within episodes of conflict.

7. Econometric Concerns

7A. Serially and Cross-Correlated Errors - Theory

Estimating response curves faces two econometric challenges particular to our approach. The data are a time-series, so we're sensitive to the possibility of *serially correlated* errors, which may reflect misspecification or omitted variables. Recall that our data are a sequence of actions, regardless of the duration of lags between actions. So any serial correlation in residuals would relate to previous residuals indexed by action (rather than calendar time), requiring consistently estimated standard errors for appropriately powered inference.

Second, and more importantly, since actions alternate between sides and sides respond to damage suffered in the previous action, *cross-correlation* of $\varepsilon^I{}_a$ with $\varepsilon^G{}_{a-1}$, implies inconsistent estimates of coefficients in (3) and (4), since $\varepsilon^G{}_{a-1}$ is correlated by construction with $D^I{}_{a-1}$ (should that residual have any variance at all).

To illustrate, consider the Israeli error term for action a, $\varepsilon^I{}_a$. Substituting (5) into (3) reveals the potential for ρ^{IG} to cause omitted variable bias

$$D^G{}_a = \beta_0 + \beta_1 D^I{}_{a-1} + \beta_2 (D^I{}_{a-1})^2 + \rho^{IG} \varepsilon^G{}_{a-1} + \rho^I \varepsilon^I{}_{a-2} + u^I{}_a \text{ for Israel, } \qquad (3').$$

ρ^{IG} nonzero implies inconsistent estimates of β_1 and β_2 (and β_0), since $Cov(D^I_{a-1}, \varepsilon^G_{a-1}) > 0$ if the regression $R^2 < 1$ in (4) (for action a-1), i.e., if $V(\varepsilon^G) > 0$, which is always the case empirically.

We can solve that omitted variable bias by adding lagged damage D^G_{a-2} to our estimating equation, but the solution introduces a new bias. Substituting from (4) (at a-1) into (3')

$$D^G_a = \beta_0 + \beta_1 D^I_{a-1} + \beta_2 (D^I_{a-1})^2 + \rho^{IG} [D^I_{a-1} - \beta_3 - \beta_4 D^G_{a-2} - \beta_5 (D^G_{a-2})^2] + \rho^I \varepsilon^I_{a-2} + u^I_a \quad (3'')$$

Now a nonzero ρ^I implies inconsistent estimates of β_4 and β_5 (and β_3), since $Cov(D^G_{a-2}, \varepsilon^G_{a-2}) > 0$ if the regression R^2 in (4) is lower than 1!

To treat both omitted variable biases at once requires adding an additional lag from the action sequence. Substituting from (3) (for action a-2) into (3'') yields:

$$D^G_a = \beta_0 + \beta_1 D^I_{a-1} + \beta_2 (D^I_{a-1})^2 + \rho^{IG} [D^I_{a-1} - \beta_3 - \beta_4 D^G_{a-2} - \beta_5 (D^G_{a-2})^2]$$
$$+ \rho^A \{D^G_{a-2} - \beta_0 - \beta_1 D^I_{a-3} + \beta_2 (D^I_{a-3})^2 \} + u^I_a \quad (7)$$

which adds damage suffered, D^I_{a-3}. Equation (7) has an error term uncorrelated with the damage variables on the right-hand side, so it is free of the omitted variable bias in (3') and (3'') (assuming --as we did above, that u^I and u^G are independent errors). It has seven reduced form coefficients and eight structural coefficients (six β's, two ρ's).[21]

A symmetric derivation yields an equation for damage suffered by Israel:

$$D^I_{a-1} = \beta_3 + \beta_4 D^G_{a-2} + \beta_5 \left(D^G_{a-2}\right)^2 + \rho^{GI}\left[D^G_{a-2} - \beta_0 - \beta_1 D^I_{a-3} - \beta_2 (D^I_{a-3})^2\right]$$
$$+\rho^G \left[D^I_{a-3} - \beta_3 - \beta_4 D^G_{a-4} - \beta_5 \left(D^G_{a-4}\right)^2\right] + u^G_{a-1} \quad (8)$$

That equation is similarly free of the omitted variable bias due to serially correlated errors, because the serial and cross-correlation coefficients are explicitly estimated by including lagged damage from actions a-3 and a-4. System (7) and (8) allow us to identify our ten structural coefficients (six β's, four ρ's), which we can estimate by nonlinear least squares.

One novel challenge presented by these data is to find a way to calculate standard errors and properly sized tests in the presence of possible correlation between residuals in this

[21] This approach is akin to a standard practice in time series analysis, of adding lagged covariates to an estimating equation until the serial correlation in residuals is resolved (e.g., Hamilton 2020). In our case, it also allows us to empirically identify structural coefficients, which the model requires that we think of as causal (Granger, 1969).

unconventional time series of residuals. This is particularly important to us, given the centrality of testing Hypotheses 1-5 to our objective of modeling deterrence. Our main approach will be to model this serial correlation and heteroskedasticity explicitly in estimating response curves in (7) and (8). We will also allow cluster-robust standard errors (i.e., robust to correlated residuals) within episodes for testing.[22]

7B. Estimation allowing for Serially and Cross-Correlated Errors

Panel B of Table 5 presents estimated coefficients of equations (7) and (8), which allow for cross-correlated and serially correlated errors. As in the simpler specification of Panel A, estimated response curves are both upward-sloping and concave. The Israeli response curve has a positive intercept, as before, while the Gazan intercept is negative and statistically insignificant (i.e., statistically zero), though not so negative that the resulting equilibrium is not violent. So our key findings from Section 6B remain: a unique stable equilibrium with incomplete deterrence (illustrated in Figure 9 below). Figure 9 also shows that, even though the simple specification has biased coefficient estimates (see discussion below), the effect on the location of equilibrium is quite small.

Turning to possible sources of bias, Panel B shows strong evidence of cross correlation: our estimates of ρ^{GI} and ρ^{IG} are large (relative to the coefficients) and statistically significant: -0.102 and -3.227 respectively. The behavioral implication is surprising. Both sides respond quite strongly with positive damage inflicted to the component of damage suffered (D_{a-1}) which our model predicts ($\hat{D}_{a-1}$), but have a far smaller (though still positive) response to the residual (unpredicted) component of damage, ε_{a-1}.[23] The linear response is reduced by 95% in the Gazan case (-0.102/0.107), and by 48% in the Israeli (-3.227/6.785). This dramatic differential response to expected damage as opposed to "unexpected" residual damage is consistent with the

---

[22] We have also tested for serial correlation by expanding on a method suggested by Wooldridge (2010) for Tobit estimators. Additionally, we've estimated the parameters of (7) using the resulting Tobit residuals, allowing cluster-robust residual variance. We then simulated response curves using estimates from (7) and our Tobit estimates, and compared results to estimate bias due to nonzero cross-correlations. These analyses are available from the authors upon request, along with the results of an alternative approach following MacKinnon et al. (2023). Results are robust to all these approaches, so we present the simplest.

[23] To see this for the Israeli response, write equation (4) as $D^I_{a-1} = \hat{D}^I_{a-1} + \varepsilon^G_{a-1}$, and substitute in (3') to get $D^G_a = \beta_0 + \beta_1 [\hat{D}^I_{a-1} + \varepsilon^G_{a-1}] + \beta_2[\hat{D}^I_{a-1} + \varepsilon^G_{a-1}]^2 + \rho^{IG} \varepsilon^G_{a-1} + \rho^I \varepsilon^I_{a-2} + u^I_a$. So $\partial D^G_a / \partial \varepsilon^G_{a-1} = \partial D^G_a / \partial \hat{D}^I_{a-1} + \rho^{IG}$. At $D^I_{a-1}=0$ (Israeli response intercept), $\partial D^G_a / \partial \varepsilon^G_{a-1} = 15.72 - 12.77 = 2.95$, 19% of $\partial D^G_a / \partial \hat{D}^U_{a-1} = 15.72$.

possibility that the sides treat those residuals as aberrations or measurement error. To the best of our knowledge, this is a novel finding that has no precedent in the conflict literature.

Figure 9: Simple, Markovian, and Fully Unconstrained Equilibria

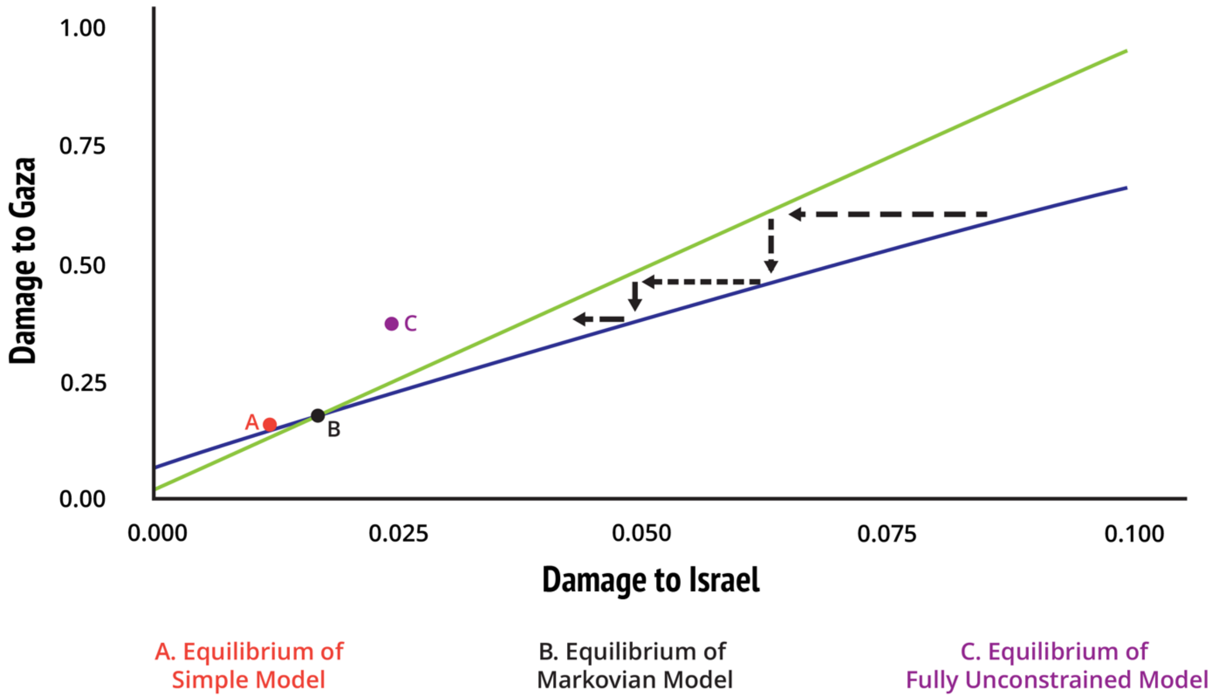

The results also show that conventional serial correlation, estimated as ρ^I and ρ^G, is present, though those estimated parameters are small and not statistically significant.

One clear message of the table is that omitting cross-correlation (and serial correlation) implies substantial omitted variable bias in the simple model (Panel A), particularly on the linear term in damage: for the Gazan side it increases threefold between columns (1) and (3), once we allow what's essentially a separate coefficient on residual damage (in equations (4) and (5)).

The large cross-correlation coefficient estimate for ρ^{IG}, of -3.227, might seem alarming given the potential implications for the variance of error terms in a dynamic system. In a conventional time series with error term $\varepsilon_t = \rho\, \varepsilon_{t-1} + u_t$, a serial correlation term $|\rho| > 1$ would imply that $V(\varepsilon_t)$ increases monotonically in t, rather than allowing convergence to some equilibrium. Our time series is different, as the sequence of damage measures alternates between sides that measure damage on different scales, so that the size of ρ^{IG} is partially due to a translation from the Gazan damage measure to the Israeli in equation (5). The dynamic implications for the variance of damage are given mainly by the *product* of the two cross-correlations ρ^{IG} x ρ^{GI}, which in our case (in panel B) is less than one (0.329 = -3.227 x -0.102). We have confirmed this dynamic stability in simulations and with a graphical exposition (not shown), and provide analytic proof for a simple case in Appendix A.

Panel B reports what will turn out to be our preferred specification. The reasoning behind that choice requires some additional specification testing and discussion of the Markov property. (This is admittedly technical; the impatient reader might be excused for skipping ahead.)

The (patient) reader might notice that we estimate coefficients ρ^{I} and ρ^{G} by including damage at lags a-3 and a-4 (i.e., for both Gaza and Israel). Implicitly, that identification was gained by assuming that the response curve (action a) is affected only by the previous lag of damage suffered (a-1) and not by the previous (a-3). That assumption, that the relevant state space is fully described by most recent history, is the Markov property in games (Markov, 1954), and allows for considerable technical simplification.

7B.1 Stability within episodes

An additional way to use action lags a-3 and a-4 is to *test* the Markov property, by estimating response coefficients directly. Panel C reports the results, for both Gaza and Israel, estimating the "Fully Unconstrained" model which allows both ρ^{I} and ρ^{G} and response curve coefficients for damage and damage squared at lags a-3 and a-4. Only one of the four added coefficients on lagged damage is statistically significant (at the 10% level), and very little explanatory power is added (the RMSE shrinks slightly from 0.270 to 0.268, and the R-squared grows slightly from 0.217 to 0.227). So Occam's razor might shave those extra coefficients at lags a-3 and a-4.

At the same time, estimated coefficients on damage and damage squared (at lags a-1 and a-2) remain significant, tracing out upwards-sloping and concave response curves, so the unique, stable equilibrium is preserved. The cross-correlations remain significant and negative, as in the "Markovian" model (Panel B).

What remains is to choose between a specification that's theoretically more tractable (the Markovian in Panel B) and another with slightly better explanatory power (Panel C). One way to think about that decision is to consider what might account for serial correlation. In principle, it could reflect behavioral responses to damage at lags a-3 and a-4; or some unobserved variable which is persistent between actions –such as weather, munitions, or the local political climate; or it could reflect measurement error or misspecification. We cannot distinguish between those possibilities with these data, so we follow the empiricist tradition of labeling our ignorance as a residual and giving it an error structure, i.e., choosing the Markovian specification. That choice has the very attractive property of stable response curves between actions within episodes.

How much difference do allowing cross-correlation and non-Markovian coefficients make to the location of equilibrium? For comparison, Figure 9 reports the equilibrium implied by the Simple Model (Panel A of Table 5, the point labeled 'A'), the Markovian Model (Panel B, Table 5; the two lines and their intersection, labeled 'B'), and the Fully Unconstrained Model (Panel C of Table 5, labeled 'C'). Equilibria A and B are extremely close, despite the omitted variable bias due to cross-correlation, and equilibrium C is not far removed.

Equilibrium C requires some explanation. In general, the model in Panel C creates a four-dimensional problem, which wouldn't lend itself to the (two dimensional) graphical analysis of Figure 9. Fortunately, analysis of an *equilibrium* simplifies: At a *stable* fixed point (D^{If}, D^{Gf}), expected reaction curves $R^G(R^I(D^{If})) = D^{If}$, $R^I(R^G(D^{Gf})) = D^{Gf}$ have $(D^{If}_{a-1}, D^{Gf}_{a-2}) = (D^{If}_{a-3}, D^{Gf}_{a-4})$. So we can graph candidate stable damage combinations against each other using all the estimated lag action coefficients (as estimated in Panel C of Table 5) to solve equilibrium and test for stability. The result is two curves similar to those depicted in Figure 9. They have a single intersection point (due to concavity), which is stable (the green cuts the blue from below). That's fortunate, as it makes the exercise internally coherent.

The implied equilibrium from that (non-Markovian) model is slightly more violent than that of the Markovian (Panel B), being further from the origin, at (0.024, 0.36), as opposed to the Markovian equilibrium at (0.015, 0.16). Overall, the equilibria implied by all three models are extremely close, reassuring us in choosing the Markovian specification of Panel B.

Referring back to our list of hypotheses to test in Section 6A, stability within episodes is (marginally) technically rejected, but we adopt it as a working hypothesis anyway –as it provides an excellent approximation and the same qualitative results: a single, stable equilibrium with incomplete deterrence.

7B.2 Stability between episodes

Another dimension of the Markovian property of games in our episodic structure is the stability of response curves *between* episodes (rather than between actions within episodes). Table 6 reports our investigation of that stability, introducing characteristics of the previous episode as right-hand side variables. For comparison, Columns (1) and (4) report our preferred specification from Table (5), the "Markovian" of Panel B. Columns (2) and (5) add as a variable accumulated damage suffered in all actions in the previous episode --an aspect of history that might affect current response.[24] In both cases, the estimated coefficient is small and statistically insignificant.

Columns (3) and (6) add as covariates the length of the lull between the previous episode and the episode preceding it.[25] A long lull might indicate an implicit or explicit truce, a decline in the intercept of either side (which reflects an inclination to conduct unprovoked attacks), or some other omitted variable --such as a decline in unobserved provocations. For both the Gazan and Israeli side estimates are negative, small, and statistically significant. The negative coefficients are intuitive. Are those coefficients of relevant size? The average lull length is 368 hours; so when multiplied by the estimated coefficient the average lull would shift the Israeli response curve by about -7% of average damage by Israel. The effect size of the Gazan lull length coefficient is larger at mean lull length, shifting by -19% of average damage by Gazan militants.

The most important point is the lack of omitted variable bias incurred by omitting this information from the previous episode. Our estimated response curves remain upward-sloping and concave, with none of the estimated β coefficients changing significantly when we add information from previous episodes. Moreover, our serial correlation and cross-correlation coefficients are also largely unchanged. We infer, then, that the simple Markovian model (with serial and cross-correlation) which uses only the most recent information from the current episode, though formally rejected --since lull length *is* predictive, provides a robust and useful approximation using the preferred specification estimated in Table 5, panel B above.

In summary, the workhorse Markov model of repeated games provides an excellent approximation of the behavior of both sides. Analysis based on damage sustained from the last action alone predicts quite well the magnitude of responses, yielding a unique and stable equilibrium in a model predicting escalation or de-escalation. Moreover, this model is tractable, provided that cross-correlation and serial correlation are accommodated in the time series.

---

[24] The patient reader will note that the estimated coefficients have changed slightly due to the smaller sample: to estimate coefficients on lagged episodes we lose eight episodes of data (two at the beginning of the sample and an additional two following each break in the sample due to a major operation).

[25] The lull length preceding the *current* episodes is endogenous by construction, since it ends with the action that begins the current episode. Adding this variable necessitates losing eight additional lagged episodes in the sample.

With a tested model in hand, we turn to the analysis of four major events: the introduction of the Iron Dome missile defense system, and the three Israeli operations in Gaza.

8. Iron Dome Missile Defense and Major Operations

We turn now to interventions that could shift response curves: the introduction of the Iron Dome anti-missile defense system, and the three major Israeli operations in our sample period, which are characterized by periods of high-frequency attacks accompanied by ground incursions in Gaza. Those major operations are not escalations in the sense of our model since they are not the result of a sequence of escalating actions (as in Figure 2) but instead pre-planned offensive operations that begin with a large violent action. During these periods the objective of the IDF, in its description, is to destroy some capacity of the other side to use violence --to deter by denial, as discussed in Section 2B above.

As illustrated in Figure 4, deterrence by denial shifts an opponent's response curve to a less aggressive stance (i.e., a smaller intercept and perhaps less slope) thereby achieving a less violent equilibrium. So, revisiting the fifth hypothesis of Section 6A, we can test whether major operations achieve deterrence by denial: does the response curve of Gazan militants shift left?

Following the logic of our model, the Iron Dome missile defense system should have the same effect. From a Gazan militant perspective, this system reduces damage imposed on Israel per attack, at the same marginal cost of attack (e.g., firing a mortar or rocket). A lower net benefit of attack should reduce the expected response at all levels of damage suffered.

Before turning to results, it's worth distinguishing temporary from persistent effects. Incursions destroy capacity (munitions and personnel), which can subsequently be restored, so we expect their effect to be temporary. Missile defense reduces the effectiveness of rockets for the remainder of the sample (permanently for the purposes of this study). That distinction informs the way we go about estimating effects. We use all the data before and after the Iron Dome missile defense system introduction to generate a comparison, since missile defense represents a permanent reduction in net benefits to Gazan militants. In contrast, for incursions, we will estimate using a six-month window before and after, since the effect on net benefits may well be temporary.

Table 7 reports the result of investigating the Iron Dome system. Columns (1) report the baseline specification from Table 5 Panel B (the Markovian), for comparison. Columns (2) augment that specification by allowing the Gazan intercept to shift when Iron Dome is introduced, in April 2011. The Gazan coefficient on Iron Dome is negative and not statistically significant. Moreover, other estimated coefficients are very inconsistent with baseline estimates (in column (1)), suggesting that we may be reaching the limit of what these data can reveal. (In particular, the linear coefficient on damage for Israel shifts to a large negative number. Note, though, that the sum of that coefficient and the Israel cross-correlation term remain about the same, suggesting that our ability to distinguish between those two coefficients is extremely sensitive to specification, because the variation available to identify the one is so multicollinear with the variation we use to identify the other.)

Columns (3) examine a more flexible specification –as suggested in the discussion of Figure 4 in Section 2, in which both Gazan *and Israeli* response curves may shift after Iron Dome is introduced. That result returns a statistically significant shift in the Gazan response curve, which is negative, as predicted by theory: a reduction in response at all levels of damage suffered. The Israeli response curve shifts in a positive direction, also as predicted by theory, though that shift is not statistically significant. Note also that the estimated linear and cross-correlation coefficients are closer to those we've seen in baseline specifications (though standard errors are much larger than in the baseline, again suggesting multicollinearity).

Figure 10 illustrates that test graphically. The solid lines illustrate response curves before the deployment of the Iron Dome system (blue for Israel, green for Gazan militants), which intersect in an incomplete deterrence equilibrium, as in Figures 8 and 9. The serrated lines illustrate response curves after Iron Dome: the Gazan response shifts left, to a less aggressive stance, as theory predicts; the Israeli shifts up, to a more aggressive stance, also as theory predicts. The Gazan shift is statistically significant though the Israeli shift is not. The net result is a lower violence equilibrium, where the two serrated lines intersect (left and below the pre-Iron Dome location). Specifically, Iron Dome missile defense shifts the equilibrium from (0.024, 0.31), to (0.004, 0.25); both sides gain by suffering less violence in the new equilibrium.

Figure 10: Does Missile Defense Deter?

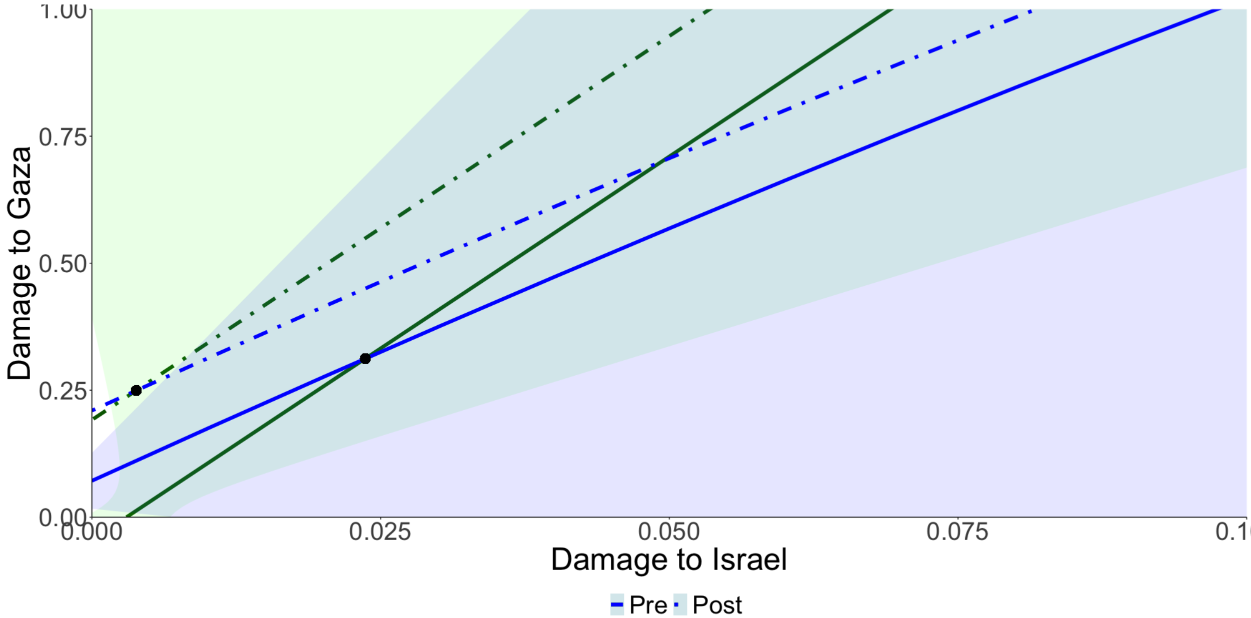

As a robustness check, columns (4) report the result of including trend terms, to allow for the possibility that the shift associated with the introduction of Iron Dome might be conflated with a trend in either the Gazan or Israeli intercept. The results are robust: the Gazan coefficient remains statistically significant and negative, while the Israeli remains positive though not statistically significant. The trend coefficients are themselves statistically insignificant. In summary, the data indicate a fairly robust finding: missile defense succeeded in shifting the Gazan response curve unequivocally to a *less aggressive* stance, while possibly shifting the Israeli response curve to a more aggressive stance.

In principle, we could apply an analogous approach to investigate the effects of the three major Israeli operations on response curves, though in practice that exercise runs out of precision. Table 8 reports the results of that analysis. For comparison, columns (1) report the estimated response curve coefficients including the Iron Dome indicator (from Table 7, column (3), estimated jointly as before). Columns (2) add estimated effects of each of the major operations, using indicator variables for each operation interacted with an indicator for the period six months before and after each. Unfortunately, this analysis seems to run into the same multicollinearity issues as did Table 7 column (2), when we attempted to allow a Gazan Iron Dome effect without an Israeli: the estimated Israeli linear damage coefficient and cross-correlation coefficient grow in opposite directions (though their sum remains approximately the same across specifications). Table 8 reports a set of estimated effects for the three major operations as a matter of due diligence, with and without trends (in columns (2) and (3) respectively) though we are reluctant to interpret them given the evidence of multicollinearity – specifically in the unstable Israeli linear coefficient. Our main conclusion from Table 8 is that we lack sufficient variation in these data to extract any further inference from our two time-series of damage.

Our major conclusion about interventions nevertheless remains: the response curve approach can be a valuable tool in analyzing shifts in curves due to interventions such as the introduction of the Iron Dome missile defense. Those shifts in response curves, in turn, imply shifts in equilibria in the context of our model.

9. Conclusions

We study deterrence in the context of the Israel-Gaza conflict, estimating response curves in a repeated game theoretical framework. Estimation is enabled by a uniquely detailed dataset based on daily reports compiled by United Nations peacekeepers, covering the period 2007 through 2017. The conflict is characterized by frequent, short episodes of violence. We code 1,170 episodes, of which the median lasts less than 5 hours, but the median pause between episodes is only about 12 days. Those episodes include almost 8,000 Israeli actions and a similar number of Gazan actions. On average, 0.11 Gazans are killed per Israeli attack and 0.41 injured, while those figures for Israelis are much lower: 0.003 and 0.03.

Our analysis reveals that this low level of persistent violence can be understood as the equilibrium of a repeated game with mixed strategies, which we solve by estimating expected response curves for each side and solving their equilibrium. That equilibrium is unique and stable: sequences of attacks and counterattacks tend to converge to it, so that even episodes beginning with large attacks will de-escalate. Yet that equilibrium is stochastic, pulling sequences toward it in probability but not holding them at a single point.

Both response curves are upward-sloping, indicating a deterrent posture. Yet the equilibrium achieved displays *incomplete deterrence* in the sense that it remains violent (i.e., it is not at the origin). This characterization is novel and relevant, we believe, to many other security situations in which deterrence is incomplete: cyber-warfare, espionage, disinformation campaigns, and other grey zone conflict, including conflict between great powers such as freedom of navigation operations at sea.

Though the data reveal that episodes are not fully independent of each other, that dependence is very small, so the textbook Markovian repeated game paradigm provides an excellent approximation of these data, describing the conflict as the combatants do. That description contrasts with the empirical literature on this conflict, which may overstate the duration (and damage) due to sequences of actions set off by an individual attack. Those studies were set in a previous period and included the West Bank, which may explain the contrast. Yet, we suspect that they also reflect a bias of the VAR approach due to aggregation of actions within episodes and conflation of episodes with each other.

We illustrate an application of the model by estimating the deterrent effect of Israel introducing a missile defense system, which shifts the Gazan response curve to a less aggressive posture, consistent with our theoretical prediction.

Tragically, the particular equilibrium we estimate for the period 2007-2017 did not survive the attack of October 7, 2023. How does the model inform our thinking about the conflict *since* those attacks? First, the current war extrapolates not only outside our sample period but represents an order of magnitude higher damage (currently about 50 times as lethal as the entire sample period for Gazans and about 75 times as lethal for Israelis). Second, the sides are no longer in a mutual deterrence equilibrium following the decimation of the fighting capacity of Hamas and other Gazan militants, so that Israel is unlikely to be deterred from operating in Gaza by the possibility of militant attacks on Israeli civilians.

Most importantly, we believe that the model provides a useful way to think about mutual deterrence in Israel's other conflicts, with the Iranian Revolutionary Guard in Iran as well as with Iranian proxies: primarily Hezbollah in Lebanon, but also Shia militants in Iraq, and Houthi forces in Yemen. Whether those conflicts are characterized by unique, stable equilibria is an empirical question of some urgency, at this writing.

We hope that our major innovation, estimating response curves within episodes of conflict, can illuminate analyses of potential escalation and de-escalation in other conflicts. It holds particular potential for policy-relevant studies of deterrence in "gray zone" conflict, such as cybersecurity and disinformation, where frequent, repeated interaction is common. Our empirical results are, of course, particular to a conflict and period.

REFERENCES

Abrahams, Alexei and Berman, Eli and Khadka, Prabin and Klor, Esteban F. and Powell, John. 2019. "Mostly Deterred: An Episodic Analysis of The Israel-Gaza Conflict" (October). http://dx.doi.org/10.2139/ssrn.3465438

Asali, Muhammad, Aamer S. Abu-Qarn, and Michael Beenstock. 2017. "The Cycle of Violence in the Second Intifada: Causality in Nonlinear Vector Autoregressive Models," *Journal of Applied Econometrics*, 32 (6): 1197–1205.

Axelrod, Robert. 1984. *The Evolution of Cooperation.* New York: Basic Books.

Baliga, Sandeep and Tomas Sjöström. 2009. "Contracting with Third Parties," *American Economic Journal: Microeconomics*, 1 (1): 75–100.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75 (5): 1331–1370.

Benmelech, Efraim, Claude Berrebi and Esteban F. Klor. 2015. "Counter-Suicide-Terrorism: Evidence from House Demolitions," *Journal of Politics*, 77 (1): 27–43.

Bergman, Ronen and Adam Goldman, "Israel Knew Hamas's Attack Plan More than a Year Ago," *New York Times*, December 2, 2023, https://www.nytimes.com/2023/11/30/world/middleeast/israel-hamas-attack-intelligence.html. (Downloaded October 8, 2024.)

Berman, Eli and David D. Laitin. 2008. "Religion, Terrorism and Public Goods: Testing the Club Model," *Journal of Public Economics*, 92 (10–11): 1942–1967.

Berman, Eli, David A. Lake, Gerard Padró I Miquel and Pierre Yared. 2019. "Principals, Agents, and Indirect Foreign Policies," in *Proxy Wars: Suppressing Violence Through Local Agents*, Eli Berman and David A. Lake (eds.). Ithaca, NY: Cornell University Press.

Chassang Sylvain and Gerard Padró i Miquel. 2010. "Conflict and Deterrence Under Strategic Risk," *The Quarterly Journal of Economics*, 125 (4): 1821–1858.

Dixit, Avinash, and Susan Skeath. 1999. *Games and Strategy.* New York: W.W. Norton & Co.

Dugan, Laura and Erica Chenoweth. 2012. "Moving Beyond Deterrence: The Effectiveness of Raising the Expected Utility of Abstaining from Terrorism in Israel," *American Sociological Review*, 77 (4): 597–624.

Durante, Ruben and Ekaterina Zhuravskaya. 2018. "Attack When the World is Not Watching? International Media and the Israeli-Palestinian Conflict," *Journal of Political Economy*, 126 (3): 1085–1133.

Etkes, Haggay and Assaf Zimring. 2015. "When Trade Stops: Lessons from the Gaza Blockade 2007-2010," *Journal of International Economics*, 95 (1): 16–27.

Evron, Yair. 2005. "Disengagement and Israeli Deterrence," *Strategic Assessment*, 8 (2): 6–11.

Fearon, James D. 1994. "Domestic Political Audiences and the Escalation of International Disputes," *The American Political Science Review*, 88 (3): 577–592.

Fearon, James D. 1997. "Signaling Foreign Policy Interests Tying Hands versus Sinking Costs," *Journal of Conflict Resolution*, 41 (1): 68–90.

Fearon, James D. 2002. "Selection Effects and Deterrence," *International Interactions*, 28 (1): 5–29.

Freedman, Lawrence. 2004. *Deterrence*. Cambridge, UK: Polity.

Fudenberg, Drew and Jean Tirole. 1991. *Game Theory*. Cambridge MA: The MIT Press.

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.

Granger, Clive W. J. 1969. "Investigating Causal Relations by Econometric Models and Cross-Spectral Methods," *Econometrica*, 37 (3): 424–438.

Gibilisco, Michael. 2023. "Mowing the Grass," *Journal of Theoretical Politics*, 35 (3): 204–231.

Golan, David and Jonathan D. Rosenblatt. 2011. "Revisiting the Statistical Analysis of the Israeli–Palestinian Conflict," *Proceedings of the National Academy of Sciences*, 108 (15): 53–54.

Hamilton, James D. 2020. *Time Series Analysis.* United States, Princeton University Press.

Haushofer, Johannes, Anat Biletzki and Nancy Kanwisher. 2010. "Both Sides Retaliate in the Israeli-Palestinian Conflict," *Proceedings of the National Academy of Sciences*, 107 (42): 17927–17932.

Heise David. 1989. Modeling Event Structures. *Journal of Mathematical Sociology*, 14 (2-3):139–69.

Hila'I, Sivan, 2023. "Netanyahu in 2017: Hamas Planned an Attack with Missiles, Hang-gliders, Kidnapping and Murder in Civilian Settlements" *Ynet*, December 12, 2023. https://www.ynet.co.il/news/article/ryd2lfgpp (downloaded October 8, 2024).

Hubbard, Ben and Maria Abi-Habib, 2023. "Behind Hamas's Bloody Gambit to Create a 'Permanent' State of War," *New York Times*, Nov. 8, 2023, https://www.nytimes.com/2023/11/08/world/middleeast/hamas-israel-gaza-war.html (downloaded October 8, 2024).

Inbar, Efraim and Eitan Shamir. 2014. "Mowing the Grass: Israel's Strategy for Protracted Intractable Conflict," *Journal of Strategic Studies*, 37 (1): 65–90.

Jaeger, David A., Esteban F. Klor, Sami H. Miaari and M. Daniele Paserman. 2012. "The Struggle for Palestinians Hearts and Minds: Violence and Public Opinion in the Second Intifada," *Journal of Public Economics*, 96 (3-4): 354–368.

Jaeger, David A., Esteban F. Klor, Sami H. Miaari and M. Daniele Paserman. 2015. "Can Militants Use Violence to Win Public Support? Evidence from the Second Intifada," *Journal of Conflict Resolution*, 59 (3): 528–549.

Jaeger, David A. and M. Daniele Paserman. 2006. "Israel, the Palestinian Factions and the Cycle of Violence," *The American Economic Review*, 96 (2): 45–49.

Jaeger, David A. and M. Daniele Paserman. 2008. "The Cycle of Violence? An Empirical Analysis of Fatalities in the Palestinian-Israeli Conflict," *The American Economic Review*, 98 (4): 1591–1604.

Jaeger, David A. and M. Daniele Paserman. 2009. "The Shape of Things to Come? On the Dynamics of Suicide Attacks and Targeted Killings," *The Quarterly Journal of Political Science*, 4 (4): 315–342.

Jervis, Robert, Richard N. Lebow and Janice G. Stein. 1985. *Psychology and Deterrence*. Baltimore, MD: Johns Hopkins University Press.

Kahn, H. 1960. *The Nature and Feasibility of War and Deterrence*. Rand Corporation.

Klinenberg, Daniel, Eli Berman and Esteban Klor. 2024. Timing is Everything: Estimating Strategic Response in Sequential Data. IGCC Working Paper.

Khoury, Jack, Yaniv Kubovich and Almog Ben Zikri. 2018. "15 Killed, Dozens Wounded as Thousands Gather on Israel-Gaza Border," *Ha'aretz*, March 31, 2018, https://www.haaretz.com/israel-news/2018-03-31/ty-article/.premium/five-palestinians-reportedly-killed-by-israeli-army-as-thousands-rally-for-mass-gaza-protests/0000017f-e121-d7b2-a77f-e3277abf0000 (downloaded October 8, 2024).

Kydd, Andrew and Barbara F. Walter. 2002. "Sabotaging the Peace: The Politics of Extremist Violence," *International Organization* 56 (2): 263–296.

Lichbach, Mark Irving, 1987. "Deterrence or Escalation? The Puzzle of Aggregate Studies of Repression and Dissent," *The Journal of Conflict Resolution*, 31 (2): 266–297.

Leventoğlu, Bahar, and Branislav L. Slantchev. 2007. "The Armed Peace: a Punctuated Equilibrium Theory of War," *American Journal of Political Science*, 51 (4): 755–771.

MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb. 2023. "Testing for the Appropriate Level of Clustering in Linear Regression Models," *Journal of Econometrics*, 235 (2): 2027–2056.

Markov, A. A., 1954. *Theory of Algorithms*. [Translated by Jacques J. Schorr-Kon and PST staff] Imprint Moscow, Academy of Sciences of the USSR, [Jerusalem, Israel Program for Scientific Translations, 1961].

Marlin-Bennett, Renee, Alan Rosenblatt, and Jianxin Wang. 1991. "The Visible Hand: The United States, Japan, and the Management of Trade Disputes," *International Interactions*, 17 (3):191–213.

Mearsheimer, John J. 1983. *Conventional Deterrence*. Ithaca, NY: Cornell University Press.

Moore, Will H. 1998. "Repression and Dissent: Substitution, Context, and Timing," *American Journal of Political Science*, 42 (3): 851–873

Myerson, Roger. 2008. "Leadership, Trust, and Power: Dynamic Moral Hazard in High Office." Working Paper, The University of Chicago.

Nanes, Matthew J. 2019. "Lebanon and Gaza, 1975-2017: Israel's Extremes of Interest Alignment," in *Proxy Wars: Suppressing Violence Through Local Agents*, Eli Berman and David A. Lake (eds.). Ithaca, NY: Cornell University Press.

Padró i Miquel, Gerard and Pierre Yared. 2012. "The Political Economy of Indirect Control," *The Quarterly Journal of Economics*, 127 (2): 947–1015.

Powell, R. 1990. *Nuclear Deterrence Theory, the Search for Credibility*. Cambridge MA: Harvard University Press.

Powell, Robert. 1999. *In the Shadow of Power: States and Strategies in International Politics.* Princeton. Princeton, NJ: Princeton University Press.

Powell, Robert. 2003. “Nuclear Deterrence Theory, Nuclear Proliferation, and National Missile Defense,” *International Security*, 27 (4): 86–118.
Powell, Robert. 2012. “Persistent Fighting and Shifting Power,” *American Journal of Political Science*, 56 (3): 620–637.
Rubin, Uzi. 2011. “The Missile Threat from Gaza: From Nuisance to Strategic Threat,” *The Begin – Sadat Center for Strategic Studies*, Bar Ilan University, Mideast Security and Policy Studies No. 91.
Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press,
Schelling, Thomas C. 1966. *Arms and Influence*. New Haven, CT: Yale University Press.
Sobelman, Daniel. 2017. “Learning to Deter: Deterrence Failure and Success in the Israel-Hizbollah Conflict, 2006-2016,” *International Security*, 41 (3): 151–196.
United Nations Office for the Coordination of Humanitarian Affairs (OCHA). 2016. *The Gaza Strip: The Humanitarian Impact of the Blockade*, November.
Wooldridge, Jeffrey M., 2010. *Econometric Analysis of Cross Section and Panel Data* (Second Edition), Cambridge MA: MIT Press.

Appendix A: Dynamic error terms in cross-correlated model

Consider our error structure in equations (5) and (6), with the standard serial correlation terms (ρ^I and ρ^G) assumed to be zero, in order to focus on the dynamic implications of the cross-correlation terms.

$\varepsilon^I_{a-1} = \rho^{IG}\, \varepsilon^G_{a-2} + u^I_{a-1}$, (5')

$\varepsilon^G_{a} = \rho^{GI}\, \varepsilon^I_{a-1} + u^G_{a}$, (6')

Note that we've arbitrarily assigned even actions to G, and odd to I.

Recursive substitution expresses ε^G_a as a function of lagged values of u^G and u^I:

$$
\begin{aligned}
\varepsilon^G_a &= \rho^{GI}\, \varepsilon^I_{a-1} + u^G_a \\
&= \rho^{GI} [\rho^{IG}\, \varepsilon^G_{a-2} + u^I_{a-1}] + u^G_a \\
&= \rho^{GI} [\rho^{IG} \{ \rho^{GI}\, \varepsilon^I_{a-3} + u^G_{a-2} \} + u^I_{a-1}] + u^G_a \\
&= \rho^{GI} [\rho^{IG} \{ \rho^{GI} [\rho^{IG}\, \varepsilon^G_{a-4} + u^I_{a-3}] + u^G_{a-2} \} + u^I_{a-1}] + u^G_a \\
&= (\rho^{GI} \rho^{IG})^2\, \varepsilon^G_{a-4} + \rho^{GI}(\rho^{GI} \rho^{IG})\, u^I_{a-3} + (\rho^{GI} \rho^{IG})\, u^G_{a-2} + \rho^{GI} u^I_{a-1} + u^G_a \\
&= .. ,
\end{aligned}
$$

continuing until we reach u^G_0. In general, that series can be written as the sum of even and odd terms,

$$\sum_{k\ even} (\rho^{IG}\, \rho^{GI})^{(k-1)/2}\, u^G_{a-k} + \sum_{k\ odd} \rho^{GI} (\rho^{IG}\, \rho^{GI})^{k/2}\, u^I_{a-k}$$

since $E(u^G_a) = 0$ for all a, it follows that $E(\varepsilon^G_a)=0$.

Now u^G_a and u^I_a are each assumed to be i.i.d , so $V(\varepsilon^G_a)$ can be written as

$$\sum_{k\ even} (\rho^{IG}\, \rho^{GI})^{(k-1)}\, V(u^G_a) + \sum_{k\ odd} (\rho^{GI})^2 (\rho^{IG}\, \rho^{GI})^{k}\, V(u^I_a)$$

These are two geometric series, so both will have bounded variance if $\rho^{IG}\, \rho^{GI}$ and $(\rho^{GI})^2 (\rho^{IG}\, \rho^{GI})^2$ are each less than one.

In Table 5, our largest estimate of $\rho^{IG}\, \rho^{GI}$ is (-10.57) x (-.0342) = 0.36 < 1, so $V(\varepsilon^G_a)$ will be bounded, even for an arbitrarily large action a, for estimates of this size.

A symmetrical argument applies to the properties of ε^I_a , $E(\varepsilon^I_a)$ and $V(\varepsilon^I_a)$.

TABLE 1

Summary Statistics on Actions

	Local		**Hybrid**		**Broad**	
	Israel	Gaza	Israel	Gaza	Israel	Gaza
Number of Actions	7616	7615	7978	7977	2550	2550
Number of Episodes	1585	1585	1170	1170	491	491
Proportion of Damage Caused | Damage >0	0.31	0.26	0.27	0.23	0.75	0.67
Number of Nonzero Actions	2387	1953	2124	1827	1916	1716
Damage Caused |Damage>0	0.42	0.04	0.47	0.04	0.53	0.05
	(0.54)	(0.05)	(0.69)	(0.06)	(0.73)	(0.06)
Airstrikes	0.26	-	0.25	-	0.78	-
	(1.18)	-	(1.24)	-	(2.06)	-
Shelling	0.31	-	0.29	-	0.92	-
	(2.38)	-	(2.56)	-	(4.47)	-
Incursion	0.15	-	0.14	-	0.45	-
	(0.43)	-	(0.45)	-	(0.8)	-
Small Arms Fire	0.11	0.04	0.1	0.04	0.32	0.11
	(0.39)	(0.2)	(0.45)	(0.2)	(0.79)	(0.34)
Mortars	-	0.53	-	0.5	-	1.57
	-	(2.41)	-	(2.47)	-	(4.2)
Rockets	-	0.53	-	0.5	-	1.58
	-	(1.97)	-	(1.91)	-	(3.03)
Grads	-	0.05	-	0.05	-	0.15
	-	(0.52)	-	(0.55)	-	(0.97)
Fatalities Caused	0.11	0.003	0.11	0.003	0.33	0.009
	(0.65)	(0.08)	(0.64)	(0.08)	(1.11)	(0.14)
Injuries Caused	0.44	0.04	0.41	0.03	1.31	0.1
	(2.34)	(0.83)	(2.58)	(0.82)	(4.5)	(1.44)

Values are the average munition or casualty per each side's non-zero actions. Lull periods are included in calculating the averages. All episodes are determined using the 48 hour lull rule. Local episodes are coded as confined to a single governorate. Broad episodes are coded ignoring governorates. Hybrid episodes link local episodes if the same munition is fired within an hour in multiple governorates. Damage is "predicted damage" as explained in discussion of Table 3 below.

TABLE 2
Episodes' Characteristics

	Local	Hybrid	Broad
Number of Episodes	1585	1170	491
A. Episode duration (number of days)			
Median	0.2	0.2	0.9
Mean	1.4	1.3	3.4
Max	120.8	121	130.1
Standard Deviation	4.2	4.6	8.7
B. Percent of episodes			
Started with Gaza Militant Violence	41.1	42.4	34.6
Ended with Gazan Militant Violence	47.3	47.8	49.2
With only Gazan Militant Violence	12	14	5
With only Israeli Violence	31	32	29
C. Gazan militant actions per episode			
Minimum	0	0	0
Median	1	1	1
Mean	1.2	1.6	3.5
Max	106	248	260
Standard Deviation	3.5	8.3	14.5
D. Israeli actions per episode			
Minimum	0	0	0
Median	1	1	1
Mean	1.5	1.8	3.9
Max	105	247	260
Standard Deviation	3.4	8.2	14.4
E. Days without attacks between episodes (lulls)			
Minimum	2	2	2
Median	4.9	6.6	3.3
Mean	8.4	11.7	4.3
Max	103.8	161.7	18.8
Standard Deviation	10.6	15.1	2.7

Calculations use the universe of firings, incursions, airstrikes, and projectile launches. Lull periods are included in calculating the averages. Episodes are determined using the 48 hour lull rule. Local episodes are coded as confined to a single governorate. Broad episodes are coded ignoring governorates. Hybrid episodes link local episodes if the same munition is fired within an hour in multiple governorates.

UN Data from June 15, 2007 until December 15, 2017.

TABLE 3
Predicted Damage by Type of Attack

	Predicted Damage to Gaza (D^G)				Predicted Damage to Israel (D^I)			
	Fatalities	**Fatalities + 0.2*Injuries**	**Fatalities + 0.5*Injuries**	**Fatalities + Injuries**	**Fatalities**	**Fatalities + 0.2*Injuries**	**Fatalities + 0.5*Injuries**	**Fatalities + Injuries**
Small Arms Fire	0.194***	0.665***	1.372***	2.550***	0.034***	0.050***	0.075***	0.116***
	(0.029)	(0.042)	(0.078)	(0.144)	(0.013)	(0.015)	(0.022)	(0.037)
Airstrikes	0.135***	0.217***	0.341***	0.546***				
	(0.017)	(0.026)	(0.04)	(0.065)				
Shellings	0.011***	0.019***	0.030**	0.049**				
	(0.004)	(0.007)	(0.012)	(0.02)				
Incursions	0.055***	0.084***	0.127***	0.199***				
	(0.016)	(0.024)	(0.037)	(0.059)				
Incursions x Small Arms Fire	−0.037	−0.463***	−1.102***	−2.167***				
	(0.056)	(0.075)	(0.119)	(0.202)				
Grads					0.002	0.034	0.082	0.163
					(0.001)	(0.022)	(0.056)	(0.112)
Rockets					0.006	0.013**	0.023***	0.040***
					(0.004)	(0.006)	(0.009)	(0.015)
Mortars					0.003	0.005	0.008*	0.013**
					(0.003)	(0.003)	(0.004)	(0.006)
Number of Attacks	3681	3681	3681	3681	3850	3850	3850	3850

Expected damage is calculated using UN reports as the unit of analysis, referred to as attack-level data in the text. Each column presents OLS coefficients of a regression of a proxy for damage on weapons used. Heteroskedasticity robust standard errors appear in parentheses. * indicates statistically significant at 10% level, ** statistically significant at 5% level; *** statistically significant at 1% level.

TABLE 4

Summary Statistics: Actions

Variable	Number of Actions	Mean	Std. Dev.	Median	Min	Max
Israeli Response						
Damage to Gaza	7978	0.13	0.41	0	0	8.23
Damage to Gaza>0	7978	0.27	0.44	0	0	1
Damage to Gaza |Damage>0	2124	0.47	0.69	0.22	0.02	8.23
Lag (Damage to Israel)	7978	0.01	0.03	0	0	0.54
Lag (Damage to Israel Squared)	7978	0.092	0.01	0	0	0.3
Lull Length	7978	367.22	494.08	186.75	48	3881
Damage to Israel Past Episode	7978	0.1	0.71	0.01	0	16.22
Damage to Israel Past Episode Squared	7978	0.51	9.86	0	0	263.11
Gazan Response						
Damage to Israel	7977	0.01	0.03	0	0	0.54
Damage to Israel>0	7977	0.23	0.42	0	0	1
Damage to Israel|Damage>0	1827	0.04	0.06	0.02	0	0.54
Lag (Damage to Gaza)	7977	0.12	0.41	0	0	8.23
Lag (Damage to Gaza Squared)	7977	0.18	1.66	0	0	67.81
Lull Length	7977	367.66	495.83	186.5	48	3881
Damage to Gaza Past Episode	7977	0.91	4.11	0.08	0	83.8
Damage to Gaza Past Episode Squared	7977	17.72	269.57	0.01	0	7022.62

Summary statistics exclude the three major operations. See Section 3 for details

The table provides summary statistics for actions. An action is a group of attacks uninterrupted by the other side, with less than 48 hours between attacks. A sequence of actions with less than 48 hours between their start times is an episode. See Section 4 for more details. Damage is predicted, using fatalities + 0.2* injuries, following discussion of Table 3 above. Lull length measured in hours.

TABLE 5
Response Curve Estimates

	A. Simple Model		B. Markovian Model		C. Fully Unconstrained Model	
	Gaza	Israel	Gaza	Israel	Gaza	Israel
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.00704*** (0.00147)	0.0820*** (0.00613)	-0.00224 (0.00175)	0.0616*** (0.00559)	0.00112** (0.00053)	0.0911*** (0.0134)
Damage	0.0322*** (0.00431)	5.825*** (0.751)	0.107*** (0.03480)	6.785*** (1.975)	0.0365*** (0.0121)	13.74*** (4.745)
Damage Squared	-0.00564*** (0.00117)	-13.39*** (2.219)	-0.000082 (0.00618)	-8.234*** (1.497)	-0.000098 (0.0025)	-7.205*** (1.317)
Lagged Damage					0.0221* (0.0121)	1.124 (1.313)
Lagged Damage Squared					-0.00164 (0.0034)	0.677 (1.564)
$\rho^{IG / GI}$			-0.102*** (0.0382)	-3.227* (1.739)	-0.0342*** (0.0108)	-10.57** (4.61)
$\rho^{I / G}$			-0.0114 (0.0462)	0.0224 (0.034)	0.125 (0.0818)	-0.0195 (0.0644)
Observations	15955	15955	15955	15955	15955	15955
Root MSE	0.281	0.281	0.270	0.270	0.268	0.268
R-squared	0.149	0.149	0.217	0.217	0.227	0.227

Observations are actions and damage is calculated using coefficients in Table 3, as described in note to Table 4. Coefficients for Gaza and Israel are jointly estimated for each panel, but reported in separate columns. Columns (1) and (2) estimate equations (3) and (4) in Section 6 by ordinary least squares. Columns (3) and (4) add own and cross-correlations as in equations (5) and (6) and are estimated as equations (7) and (8) of Section 7 by nonlinear least squares. Lagged damage refers to previous actions within episode, augmenting equations (7) and (8), also estimated by nonlinear least squares.

Standard errors, clustered by episode, are heteroskedasticity-robust and appear in parentheses. * indicates statistically significant at 10% level, ** statistically significant at 5% level; *** statistically significant at 1% level.

TABLE 6

Testing Stability Across Episodes

	Israel			Gaza		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.0616***	0.0650***	0.0942***	-0.00224	-0.00100	0.00104
	(0.00559)	(0.00914)	(0.0107)	(0.00175)	(0.00237)	(0.000930)
Damage	6.785***	8.364*	8.288***	0.107***	0.0822	0.0709**
	(1.975)	(4.582)	(2.928)	(0.0348)	(0.0518)	(0.0311)
Damage squared	-8.234***	-8.208***	-7.939***	-0.0000815	0.0000185	0.000240
	(1.497)	(1.515)	(1.489)	(0.00618)	(0.00455)	(0.00444)
$\rho^{IG / GI}$	-3.227*	-4.812	-4.847*	-0.102***	-0.0770*	-0.0672***
	(1.739)	(4.285)	(2.698)	(0.0382)	(0.0435)	(0.0237)
$\rho^{I / G}$	-0.0114	-0.00875	-0.00972	0.0224	-0.00327	0.0197
	(0.0462)	(0.0346)	(0.0336)	(0.0335)	(0.0682)	(0.0503)
Damage Previous Episode		-0.00284	-0.00464		0.000593	0.000525
		(0.00382)	(0.00379)		(0.000489)	(0.000503)
Lull Length			-0.0000243***			-0.00000516**
			(0.00000828)			(0.00000247)
Observations	15955	15755	15738	15955	15755	15738
R-squared	0.217	0.217	0.219	0.217	0.217	0.219
Pr(Previous Damage=0 & Lull Length=0)			0.012			0.0086

Observations are actions, and damage is calculated using coefficients in Table 3 as weights, as described in note to Table 4. Coefficients for Gaza and Israel are jointly estimated for each specification (as in Table 5), but reported in separate columns. For instance, coefficient estimates reported in columns (1) and (4) are of equations (7) and (8) in Section 7, jointly estimated by nonlinear least squares. Coefficient estimates reported in columns (2) and (5), and those repoted in columns (3) and (6) are also jointly estimated. Damage in the previous episode is the sum of all damage suffered by the side in the column header in the previous episode. Lull length is between the previous episode and the episode preceding it.

Standard errors, clustered by episode, are heteroskedasticity-robust and appear in parentheses. * indicates statistically significant at 10% level, ** indicates statistically significant at 5% level; *** indicates statistically significant at 1% level.

TABLE 7
The Effects of Iron Dome on Response Curves

	Israel				Gaza			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Intercept	0.0616***	-0.233	0.0712**	0.0716**	-0.00224	0.00640*	0.00368*	0.00299
	(0.00559)	(0.495)	(0.0295)	(0.0280)	(0.00175)	(0.00361)	(0.00205)	(0.00192)
Damage	6.785***	-40.98	9.553	10.37	0.107***	0.000994	0.0718*	0.0663*
	(1.975)	(56.87)	(5.874)	(6.731)	(0.0348)	(0.0133)	(0.0414)	(0.0394)
Damage Squared	-8.234***	-8.523***	-8.569***	-8.561***	-0.0000815	-0.00000595	-0.0000891	-0.0000784
	(1.497)	(1.575)	(1.563)	(1.549)	(0.00618)	(0.000534)	(0.00387)	(0.00343)
$\rho^{IG/GI}$	-3.227*	44.69	-5.851	-6.662	-0.102***	0.00443	-0.0656	-0.0598
	(1.739)	(56.65)	(5.854)	(6.698)	(0.0382)	(0.0129)	(0.0410)	(0.0381)
$\rho^{I/G}$	-0.0114	0.00193	-0.0120	-0.0108	0.0224	0.217	-0.0150	-0.0255
	(0.0462)	(0.00487)	(0.0293)	(0.0263)	(0.0335)	(0.318)	(0.0921)	(0.0966)
Iron Dome			0.112	0.138		-0.000647	-0.0140***	-0.0157***
			(0.106)	(0.141)		(0.000908)	(0.00464)	(0.00526)
Trend				-0.0000418				0.00000536
				(0.0000746)				(0.00000533)
Observations	15955	15955	15955	15955	15955	15955	15955	15955
R-squared	0.217	0.218	0.218	0.218	0.217	0.218	0.218	0.218
RMSE	0.270	0.270	0.270	0.270	0.270	0.270	0.270	0.270

Observations are actions and damage is calculated using coefficients in Table 3 as weights, as described in note to Table 4. Coefficients for Gaza and Israel are jointly estimated for each column number, but reported in separate columns (as in Tables 5 and 6). Coefficient estimates are of equations (7) and (8) in Section 7, including own and cross-correlations as in equations (5) and (6), and are estimated by nonlinear least squares. The Iron Dome variable is an indicator which begins taking the value one in April 2011 when the Iron Dome batteries were activated.

Standard errors, clustered by episode, are heteroskedasticity-robust and appear in parentheses. * indicates statistically significant at 10% level, ** indicates statistically significant at 5% level; *** indicates statistically significant at 1% level.

TABLE 8

The Effects of Iron Dome and Major Operations on the Response Curves

	Israel			**Gaza**		
	(1)	(2)	(3)	(1)	(2)	(3)
Intercept	0.0712**	0.117**	0.0957*	0.00368*	0.00188*	0.00115*
	(0.0295)	(0.0589)	(0.0504)	(0.00205)	(0.00103)	(0.000653)
Damage	9.553	57.25***	73.08***	0.0718*	0.0175**	0.0143**
	(5.874)	(17.22)	(23.99)	(0.0414)	(0.00752)	(0.00656)
Damage Squared	-8.569***	-9.120***	-9.140***	-0.0000891	-0.0000215	-0.0000162
	(1.563)	(1.598)	(1.597)	(0.00387)	(0.000455)	(0.000351)
$\rho^{IG / GI}$	-5.851	-53.24***	-69.06***	-0.0656	-0.0125***	-0.0103***
	(5.854)	(17.04)	(23.84)	(0.0410)	(0.00441)	(0.00396)
$\rho^{I / G}$	-0.0120	-0.0038	-0.00291	-0.0150	-0.271***	-0.300***
	(0.0293)	(0.00406)	(0.00322)	(0.0921)	(0.0806)	(0.0821)
Iron Dome	0.112	-0.239	-0.175	-0.0140***	0.00256	0.00113
	(0.106)	(0.479)	(0.441)	(0.00464)	(0.00651)	(0.00444)
Cast Lead		-0.209	-0.189		0.00447	0.00337
		(0.444)	(0.416)		(0.00601)	(0.00433)
Pillar of Defense		1.597**	1.609**		-0.0208**	-0.0152**
		(0.814)	(0.798)		(0.00877)	(0.00691)
Protective Edge		0.0492	0.0909		-0.0000174	-0.000250
		(0.313)	(0.285)		(0.00369)	(0.00244)
Trend			-0.00157**			0.0000161**
			(0.000740)			(0.0000073)
Observations	15955	15955	15955	15955	15955	15955
R-squared	0.218	0.233	0.234	0.218	0.233	0.234
RMSE	0.270	0.267	0.267	0.270	0.267	0.267

Observations are actions, and damage is calculated using coefficients in Table 3 as weights, as described in note to Table 4. Coefficients for Gaza and Israel are jointly estimated for each column number, but reported in separate columns. Coefficient estimates are of equations (7) and (8) in Section 7, including own and cross-correlations as in equations (5) and (6), and are estimated by nonlinear least squares. The Iron Dome variable is an indicator which begins taking the value one in April 2011 when the Iron Dome batteries were activated. Cast Lead, Pillar of Defense and Protective Edge operations are similarly measured by an indicator indicating the period after the operation. An additional indicator allowing for an intercept shift (for Israel and Gaza) is included to indicate the period six months before and after each of those operations. The coefficient on that indicator is not reported in the

Standard errors, clustered by episode, are heteroskedasticity-robust and appear in parentheses. * indicates statistically significant at 10% level, ** statistically significant at 5% level; *** statistically significant at 1% level.