INSTITUTE for REPLICATION

No. 89
I4R DISCUSSION PAPER SERIES

Comment on "Gone with the Wind: The Consequences of US Drone Strikes in Pakistan"

Joanne Haddad
Moritz Hennicke
Lamis Kattan
Thibault Lemaire
Marco Musumeci

November 2023

INSTITUTE for REPLICATION

I4R DISCUSSION PAPER SERIES

I4R DP No. 89

Comment on "Gone with the Wind: The Consequences of US Drone Strikes in Pakistan"

Joanne Haddad[1], Moritz Hennicke[1], Lamis Kattan[2], Thibault Lemaire[1], Marco Musumeci[3]

[1]Université libre de Bruxelles/Belgium
[2]Georgtown University of Qatar
[3]Erasmus University Rotterdam/The Netherlands

NOVEMBER 2023

E-Mail: joerg.peters@rwi-essen.de
RWI – Leibniz Institute for Economic Research
Hohenzollernstraße 1-3
45128 Essen/Germany
www.i4replication.org

ISSN: 2752-1931

Comment on "Gone with the Wind: The Consequences of US Drone Strikes in Pakistan"*

Joanne Haddad† *Université libre de Bruxelles*
Moritz Hennicke ‡ *Université libre de Bruxelles*
Lamis Kattan§ *Georgetown University in Qatar*
Thibault Lemaire ¶ *Université libre de Bruxelles*
Marco Musumeci|| *Erasmus University Rotterdam*

November, 2023

Abstract

Mahmood and Jetter (2023) rely on daily wind conditions as an exogenous source of variation to assess the effects of 420 US drone strikes conducted in Pakistan from 2006 to 2016. The findings indicate that these drone strikes promote a subsequent surge in terrorism over the following days and weeks, contributing significantly to as much as 19% of all terrorist incidents and resulting in over 3,000 casualties in Pakistan during the specified period. In this comment, we successfully reproduce all the results from Mahmood and Jetter (2023), including tables and figures. We then conduct four sensitivity analyses to confirm the primary findings outlined in the original paper. We document the robustness of the main results in three out of four sensitivity checks, involving the omission of all controls across various specifications, utilization of the *fixest* package in *R*, and the inclusion of control variables determined through Lasso regressions. However, we show that the addition of year fixed effects substantially reduces the first-stage F-statistics and challenges the established negative relationship between wind gusts and drone strikes.

*We thank Abel Brodeur for comments and suggestions. All remaining errors are our own.
†European Center for Advanced Research in Economics and Statistics (ECARES). Email: joanne.haddad@ulb.be.
‡European Center for Advanced Research in Economics and Statistics (ECARES). Email: moritz.hennicke@ulb.be.
§Edmund A. Walsh School of Foreign Service. Email: lk811@georgetown.edu.
¶European Center for Advanced Research in Economics and Statistics (ECARES). Email: thibault.lemaire@ulb.be.
|| Erasmus School of Economics. Email: musumeci@ese.eur.nl.

Introduction

To address the challenges of studying drone strikes' impact, Mahmood and Jetter (2023) introduces an innovative identification strategy that uses wind conditions as an instrumental variable. Contrary to prior correlational studies, their findings using this method suggest that drone strikes actually lead to an increase in subsequent terror attacks, affecting both the timing and overall number of attacks and deaths. Additionally, the paper examines the public's emotional responses to drone strikes in Pakistan, revealing that these actions trigger negative emotions and increased anger towards the United States in media coverage. Furthermore, there is a rise in anti-U.S. protests, while protests against terrorist groups remain relatively unchanged. This paper's findings have broad implications for foreign military interventions, counter-terrorism efforts, understanding the drivers of terrorism, and assessing the impact of U.S. actions on anti-U.S. sentiment and radicalization.

The empirical strategy relies on a unique dataset spanning from 2006 to 2016, focusing on drone strikes in Pakistan. To address the challenge of causally attributing effects to drone strikes, the authors employ an innovative identification strategy. They use wind conditions as an instrumental variable (IV), hypothesizing that drone strikes are less likely to occur on windy days due to operational limitations. This IV helps isolate the causal impact of drone strikes from other factors. The dataset also allows the examination of various specifications, alternative IV definitions, econometric methods, and control variables to ensure robustness in their findings. This comprehensive empirical approach offers a rigorous assessment of the relationship between drone strikes, subsequent terrorism, and public sentiment.

In this comment, we provide an overview from reproducing and replicating Mahmood and Jetter (2023). As a first step, we successfully reproduce all the results in the paper, including all tables and figures. All results were fully reproducible and match the published version of the paper. Reproduction required additional packages, which were stated clearly in the "read me" file. We present the authors' main results (Table 2) in Table 1 of this comment.

Next, we carry out four sensitivity analyses. First, we examine how the main results change from omitting all controls from the various specifications and examine the raw relationship between drone strikes and subsequent terror attacks. Second, to account for plausible year specific characteristics in both drone strikes and terror attacks, we investigate whether the main results remain robust to the inclusion of year fixed effects. Third, we repeat the authors' main results using the *fixest* package in *R*. Lastly, we repeat the main results, but with controls variables selected in Lasso regressions.

Sensitivity Analysis

Our initial assessment, presented in Table 2, affirms the primary findings outlined in the paper. We show that when wind gusts or wind speed are used as instrumental variables, drone strikes emerge as a positive and statistically significant predictor of subsequent terror attacks, even when considering the raw relationship between drown strikes and subsequent terror attacks. The remaining columns of Table 2 demonstrate that consistent results are obtained when three-, six-, or 14-day periods as the units of observation are used. If anything, the exclusion of controls from the analysis leads to an increase in the point estimates.

Our second sensitivity analysis involves replicating the same analysis presented in the main table of Mahmood and Jetter (2023)'s paper, with the addition of year fixed effects in the model specification. We present the results in Table 3. It is important to highlight that when consid-

ering year-specific effects, there is a significant decrease in the first-stage F-statistic, leading to a larger bias in the 2SLS estimates. This indicates that year by year, the negative relationship between wind gusts and drone strikes does not hold. To illustrate this, we replicate the binned scatterplots (similar to Figure 2 (a) in the paper) and account for year fixed effects in Figure 1. The figure shows that yearly specific characteristics absorb the negative relationship between the independent variable and the instrument.

To further investigate this matter, we replicate Figure 2 (a) for each individual year from 2006 to 2016. The results are presented in Appendix Figure A1. As illustrated, the link between wind gusts and drones varies from year to year, occasionally displaying a positive correlation, such as in 2013 and 2014.

In Figure 2, we replicate Appendix Figure E1 from Mahmood and Jetter (2023). However, instead of visualizing the evolution of drone strikes and terror attacks over the entire sample period using kernel-weighted local polynomial smoothing, we focus on visualizing the variations in drone strikes and wind gusts. While the relationship is vague on a year-by-year basis, the overall relationship appears to be negative. This implies that in years with less wind, more drone strikes occurred. However, as indicated by Figure A1, the within year variation in wind gusts is not clearly linked with the within year variation in drone strikes.

Next, we repeat the authors' main results using the *fixest* packaged in *R*. Results are reported in Table 4. Estimates remain statistically significant throughout.

Lastly, we rely on Lasso regressions to select control variables using all variables from the baseline specification plus additional lags of Pakistan military action. We provide a visualization regarding the selection of control variables in Appendix Figure A2. In Table 5, we repeat the main results from Mahmood and Jetter (2023), using selected controls from Lasso regressions. We show that the results remain robust.

Conclusion

This comment shows that the main results in Mahmood and Jetter (2023) remain robust to various sensitivity analyses, encompassing the exclusion of all controls across different specifications, application of the *fixest* package in *R*, and the incorporation of control variables chosen through Lasso regressions. Notably, when replicating the primary analysis with the inclusion of year fixed effects, a substantial reduction in the first-stage F-statistic is observed, pointing to an increased bias in the 2SLS estimates. We document that adding year fixed effects to the analysis challenges the negative relationship between wind gusts and drone strikes. In fact, yearly variations from 2006 to 2016 display an inconsistent correlation, occasionally displaying a positive relationship.

References

Mahmood, R. and Jetter, M.: 2023, Gone with the wind: The consequences of us drone strikes in pakistan, *The Economic Journal* **133**(650), 787–811.

Figure 1: Wind gusts in Miran Shah and drone strikes. Adding year fixed effects.

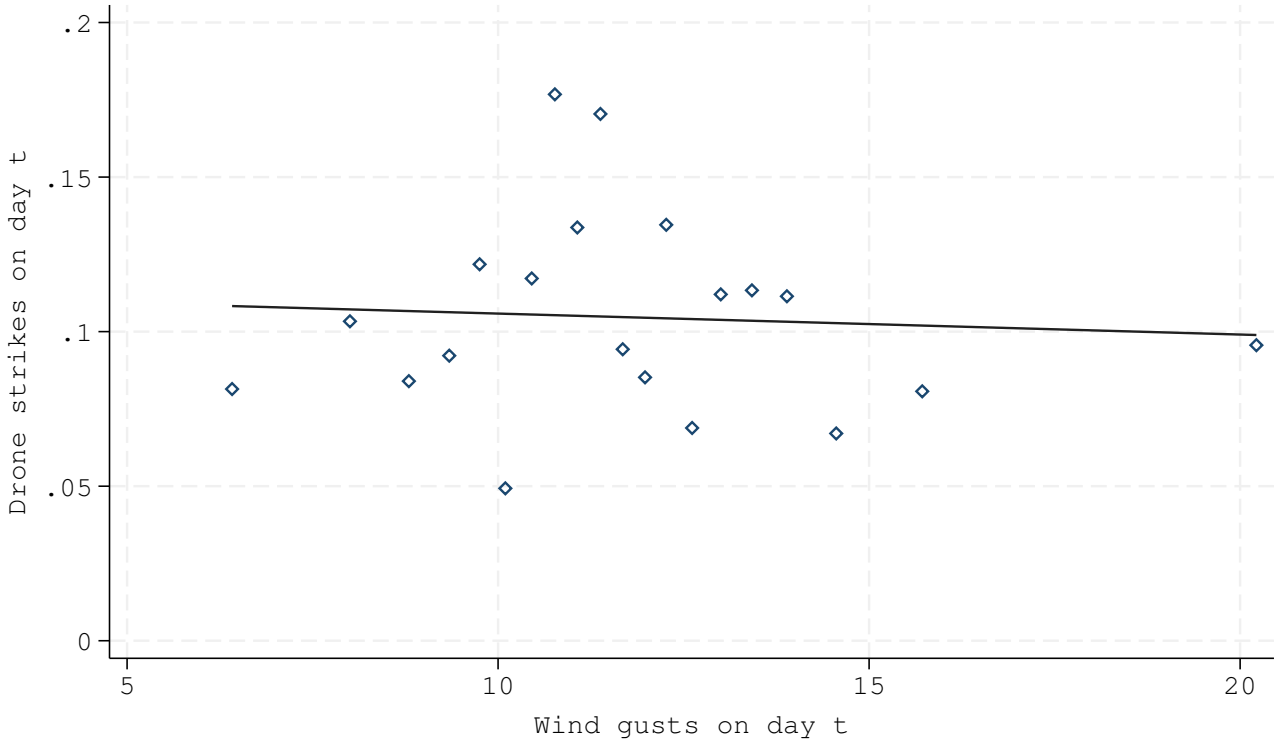

Notes: This figure replicates Figure 2(a) from Mahmood and Jetter (2023) and represents binned scatterplots of wind conditions *(x Axis)* against the number of contemporary drone strikes in the subsequent seven days at various locations. In addition to the control variables included by the authors, **we account for year fixed effects.**

Figure 2: Drone strikes and wind gusts

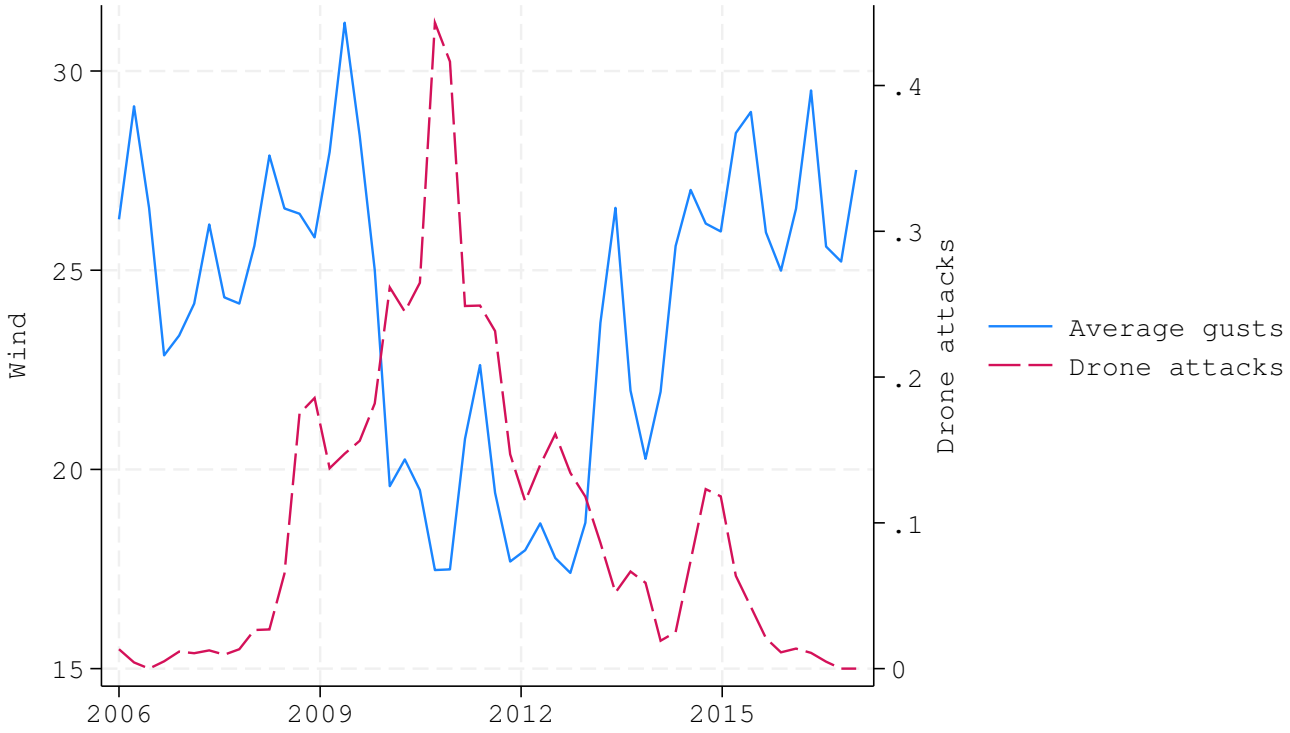

Notes: This figure replicates Appendix Figure E1 from Mahmood and Jetter (2023) by visualizing the change in drones strikes and wind gusts over the entire sample period , employing Kernel-weighted local polynomial smoothing.

Table 1: Authors' main results: All regressions account for the full set of control variables

	Days t+ 1 until 7		Three-days averages	Six-days averages	14-days averages
	(1)	(2)	(3)	(4)	(5)
			Panel A: second-stage-results, predicting subsequent terror attacks		
Drones strikes	2.9268***	3.3847**	3.0358***	3.4046**	2.5263*
	(1.0159)	(1.3543)	(1.4397)	(1.5034)	(1.4517)
	[0.004]	[0.012]	[0.035]	[0.024]	[0.082]
			Panel B: first-stage-results, predicting drone strikes on day t		
Wind gusts	-0.0028***		-0.0043***	-0.0061***	-0.0094***
	(0.0006)		(0.0009)	(0.0012)	(0.0021)
	[0.000]		[0.000]	[0.000]	[0.000]
Wind speed		-0.0054***			
		(0.0015)			
		[0.000]			
			Panel C: statistical properties		
fstat	19.388***	13.494***	21.767**	25.772**	19.188**
	[0.000]	[0.0002]	[0.0000]	[0.0000]	[0.0000]
N	3,997	3,997	1,333	666	285

Notes: This table reproduces the main results of Mahmood and Jetter (2023) reported in Table 2. All IV estimations are conducted using the *ivreg2* command in Stata with robust, heteroskedastic and autocorrelation consistent (HAC) standard errors (option *r bw(1)* in Stata) displayed in parenthesis: *** p<0.01, ** p<0.05, * p<0.1. p-values are displayed in brackets. All columns includes 14 measures of terror attacks in the preceding 14 days, 14 measures of Pakistani military actions in the preceding 14 days, a binary indicator for Ramadan, a linear time trend, measures for temperature and precipitation in Miran Shah as well as fixed effects for each day of the week and month of the year. In columns (3)–(5), we follow the authors and account for lagged values of Pakistani military actions and terror attacks for five, two and two periods, respectively. Day-of-the-week fixed effects are excluded in column (5). In columns (3)-(5), the dependent variables are measured over the following three, six and 14 days, respectively.

Table 2: Main Results, No Controls

	Days t+ 1 until 7		Three-days averages	Six-days averages	14-days averages
	(1)	(2)	(3)	(4)	(5)
		Panel A: second-stage-results, predicting subsequent terror attacks			
Drone strikes	7.5160***	21.2181***	6.5860***	6.3864***	7.6445***
	(2.2823)	(7.1628)	(1.4607)	(1.0567)	(0.8092)
	[0.000]	[0.0002]	[0.0000]	[0.0000]	[0.0000]
		Panel B: first-stage-results, predicting drone strikes on day t			
Wind gusts	-0.0025***		-0.0039***	-0.0053***	-0.0074***
	(0.0006)		(.0005)	(.0004)	(.0004)
	[0.000]		[0.0000]	[0.000]	[0.000]
Wind speed		-0.0039 ***			
		(0.0012)			
		[0.0002]			
		Panel C: statistical properties			
fstat	19.39***	13.49***	21.76**	25.72**	19.18**
	[0.000]	[0.0002]	[0.0000]	[0.0000]	[0.0000]
N	4011	4011	4013	4007	3991

Notes: This table presents the Authors' main results but excludes all controls. In all columns we **exclude** the 14 measures of terror attacks in the preceding 14 days, 14 measures of Pakistani military actions in the preceding 14 days, the binary indicator for Ramadan, the linear time trend, measures for temperature and precipitation in Miran Shah as well as the fixed effects for each day of the week and month of the year. In columns (3)–(5), we **exclude** the lagged values of Pakistani military actions and terror attacks for five, and the two and two periods, respectively. Day-of-the-week fixed effects remain excluded in column (5). All IV estimations are conducted using the *ivreg2* command in Stata with robust, heteroskedastic and autocorrelation consistent (HAC) standard errors (option *r bw(1)* in Stata) displayed in parenthesis: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. p-values are displayed in brackets. In columns (3)-(5), the dependent variables are measured over the following three, six and 14 days, respectively.

Table 3: Authors' main results: All regressions account for the full set of control variables. *Adding Year Fixed Effects*

	Days t+ 1 until 7		Three-days averages	Six-days averages	14-days averages
	(1)	(2)	(3)	(4)	(5)
			Panel A: second-stage-results, predicting subsequent terror attacks		
drones strikes	-3.7449	13.281	-7.25396	-12.9270	-15.269
	(5.111)	(27.0287)	(8.6558)	(41.529)	(1.4517)
	[0.464]	[0.623]	[0.402]	[0.446]	[0.713]
			Panel B: first-stage-results, predicting drone strikes on day t		
drones strikes	0.00062		0.000977	0 .00094	0.00071
	(0.0006)		(0.00082)	(0.00106)	(0.00170)
	[0.304]		[0.239]	[0.377]	[0.675]
wind speed		-0.000672			
		(0.00131)			
		[0.609]			
			Panel C: statistical properties		
fstat	1.06	0.26	1.39	0.78	0.18
	[0.3040]	[0.6089]	[0.2387]	[0.3769]	[0.6754]
N	3997	3997	1333	666	285

Notes: This table presents the Authors' main results. All IV estimations are conducted using the *ivreg2* command in Stata with robust, heteroskedastic and autocorrelation consistent (HAC) standard errors (option *r bw(1)* in Stata) displayed in parenthesis: *** p<0.01, ** p<0.05, * p<0.1. p-values are displayed in brackets. All columns includes 14 measures of terror attacks in the preceding 14 days, 14 measures of Pakistani military actions in the preceding 14 days, a binary indicator for Ramadan, a linear time trend, measures for temperature and precipitation in Miran Shah as well as fixed effects for each day of the week and month of the year. **We additionally include year fixed effects in all columns.** In columns (3)–(5), we follow the authors and account for lagged values of Pakistani military actions and terror attacks for five, two and two periods, respectively. Day-of-the-week fixed effects are excluded in column (5). In columns (3)-(5), the dependent variables are measured over the following three, six and 14 days, respectively.

Table 4: Authors' main results: Using the *fixest* package in *R*

	Attacks	Attacks	Attacks
Fitted Drone Strikes	2.927**	2.927***	2.927***
	(1.162)	(1.022)	(1.022)
	[0.012]	[0.004]	[0.004]
Num.Obs.	3997	3997	3997
Std.Errors	IID	Heteroskedasticity-robust	Newey-West (L=7)

Notes: This table presents the authors' main results using the *fixest* package in *R*. Standard errors are reported in round brackets and p-values in square brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Authors' main results: Selecting controls with Lasso regressions

	No Controls	No Controls (Net of Missing Controls)	Baseline Controls	Lasso Controls
Fitted Drone Strikes	7.516***	7.497***	2.927***	2.557***
	(2.283)	(2.272)	(1.022)	(0.942)
	[0.001]	[0.001]	[0.004]	[0.007]
Num.Obs.	4011	3997	3997	3963
Std.Errors	Newey-West (L=7)	Newey-West (L=7)	Newey-West (L=7)	Newey-West (L=7)

Notes: This table presents results with control variables selected in Lasso regressions using the *glmnet* package in *R*. The baseline outcome, endogenous variable and instrument are hereby regressed on a large number of controls. All controls whose coefficients were not shrunk to zero in one of the three Lasso regressions were included in the final IV regression using *fixest* in *R*. The pool of controls includes all baseline control plus lags of military action of Pakistan's military. Standard errors are reported in round brackets and p-values in square brackets. *** p<0.01, ** p<0.05, * p<0.1.

Figure A1: Year-by-Year. Wind gusts in Miran Shah and drone strikes.

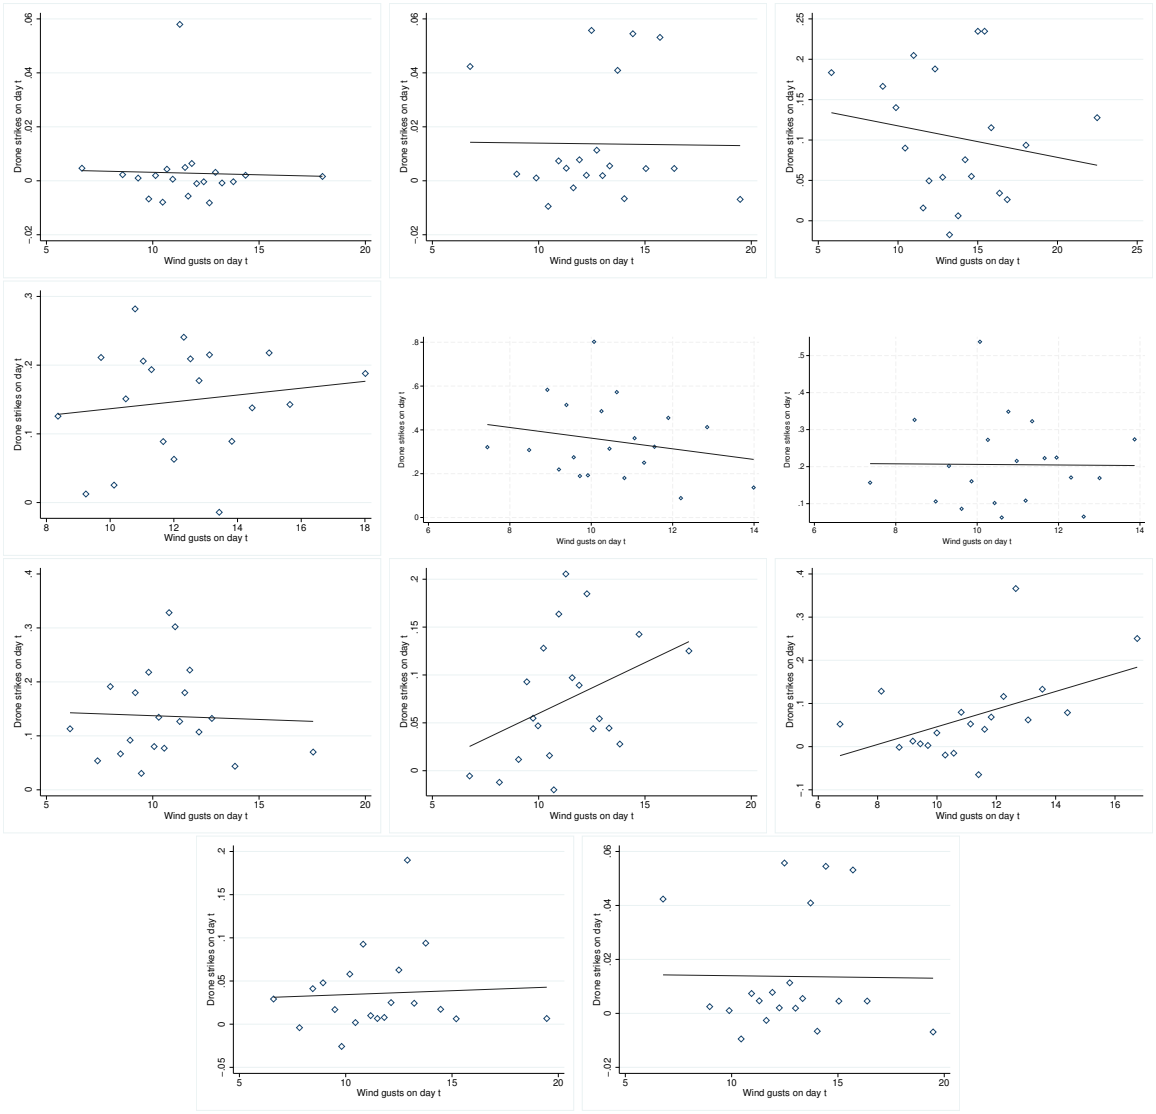

Notes: This figure replicates authors Figure 2(a) and represents binned scatterplots of wind conditions *(x Axis)* against the number of contemporary drone strikes in the subsequent seven days at various locations for ever year between 2006 and 2016.

Figure A2: Selection of controls through Lasso regression

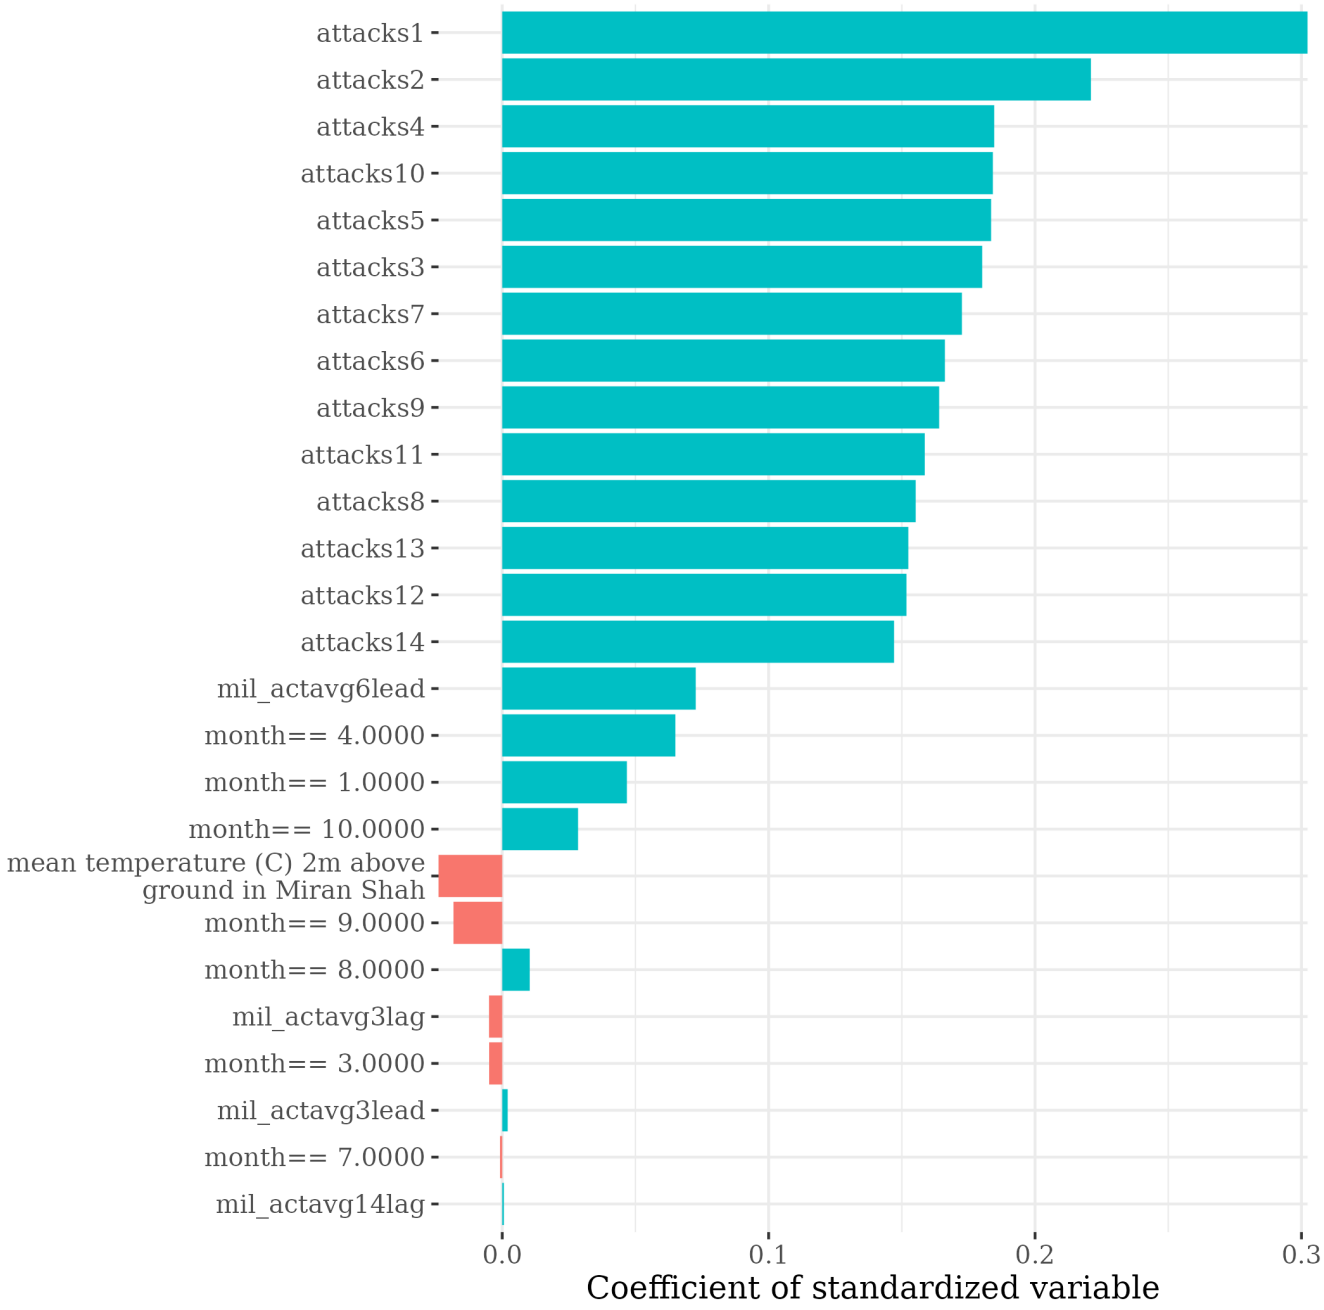

Notes: This figure visualizes the selection of control variables through Lasso regressions regressing in this case the outcome on all controls. The controls are standardized and coefficients can thus be ranked by their predictive power of variations of the outcome.