The Brides of Boko Haram: Economic Shocks, Marriage Practices, and Insurgency in Nigeria*

Jonah Rexer†

University of Pennsylvania

June 28, 2019

Abstract

Unmarried young men are often believed to cause social unrest. This paper documents that imbalances in the marriage market are indeed linked to greater violence in the context of the Boko Haram insurgency in Nigeria. Marriage markets in rural Nigeria are characterized by the customs of bride-price – pre-marital payments from the groom to the family of the bride – and polygamy. These norms diminish marriage prospects for young men, causing them to join violent insurgencies. Using an instrumental variables strategy, I find that greater inequality of brides among men increases the incidence of militant activity by Boko Haram. To instrument for marriage inequality, I exploit the fact that young women delay marriage in response to good pre-marital economic conditions, which increases marriage inequality more in polygamous villages. Supporting the mechanism, I find that the same positive female income shocks which increase marriage inequality and extremist activity also reduce female marriage hazard, lead women to marry richer and more polygamous husbands, generate higher average marriage expenditures, and increase abductions and violence against women. The results shed light on the marriage market as an important but hitherto neglected driver of violent extremism.

JEL Classifications: D74, J12, Z12, Z13

Keywords: marriage market, traditional cultural practices, extremism, conflict, income shocks

---

*For their outstanding feedback and assistance, I thank Camilo Garcia-Jimeno, Guy Grossman, and Nina Harrari. Santosh Anagol, Corrine Low, Robert Jensen, Sunny Lee, Hillary Matfess, and participants at the Wharton Business Economics and Public Policy graduate student seminar also provided helpful comments. I also thank participants at the Pacific Development Conference, the Development and Political Economics graduate student conference at Stanford, the Midwest Development Conference, the Empirical Studies of Conflict annual meeting at UCSD, and three anonymous referees, for their valuable feedback. This work has been supported by the Amy Morse grant at the University of Pennsylvania.

†Wharton School of Business. 3620 Locust Walk, Philadelphia, PA 19104. Email: jorexer@wharton.upenn.edu. `https://bepp.wharton.upenn.edu/profile/jorexer/`

1 Introduction

As youth populations swell, policymakers in many poor countries are seeking answers to a long-standing question in social science: why do young men often gravitate toward organized violence, whether in the form of rebel groups, gang violence, or religious extremism? The rational choice view argues that individuals and groups optimally trade off the returns to violence and opportunity costs of fighting. The evidence generally supports this view: economic shocks such as commodity prices fluctuations, climatic events, and employment changes that make people better off typically reduce conflict by raising opportunity costs, whereas shocks that increase the value of rebellion – particularly high prices of natural resources – will increase violence.

The marriage market can play a role in structuring the incentives for violence; the "bare branches" hypothesis holds that a surplus of unmarried men can destabilize society (Hudson and den Boer 2004). If this is true, then traditional marriage practices may be at the root of violence in countries plagued by political instability. Across sub-Saharan Africa, marriage markets are characterized by two such practices: bride price – payments from the groom to the bride on marriage – and polygamy – the practice of taking more than one wife.[1] The idea that these practices may promote violence has recently gained interest in the popular press.[2] In Northern Nigeria, the Islamist group Boko Haram has "continued to organize inexpensive weddings for its members...[who] would have probably remained bachelors" (Hudson and Matfess, 2017). These insights suggest a grim logic underlying the use of gender-based violence among these groups, in particular mass abductions of women and girls[3]: men with limited marriage prospects face low opportunity costs to joining violent groups, allowing insurgent groups to use marriage as a recruitment tool.

In this paper, I show that marriage inequality – measured as the Gini coefficient of the distribution of brides among men in a village – increases the incidence extremist violence in Nigeria. To guide the empirical exercise, I develop a model of occupational choice and marriage in which men choose between joining insurgent groups or participating in the civilian marriage market. When women experience positive income shocks, they raise their reservation bride price, pushing them toward richer mates. This dynamic is more pronounced when polygamy is allowed. Intuitively, in polygamous places, women can match with wealthier already-married men, yielding greater option value

[1] While polygamy refers both to the practice of taking many husbands (polyandry) and wives (polygyny), I will use it to mean the latter, given the relative absence of the former in Africa and the Middle East.

[2] The Economist, for example has recently run several stories with headlines such as "The link between polygamy and war" and "Why polygamy breeds civil war".

[3] In 2014-15 alone, Amnesty International estimates that Boko Haram kidnapped 2,000 girls.
https://www.amnesty.org/en/latest/news/2015/04/nigeria-abducted-women-and-girls-forced-to-join-boko-haram-attacks/

of waiting and driving up bride prices. These future payoffs also determine the sensitivity of reservation prices to an income shock. Therefore, income shocks reduce the value of civilian life more in polygamous places, pushing more men into insurgent activity. The model suggests that the interaction between female income shocks and village-level polygamy norms can be used as an instrument to address the endogeneity of marriage inequality.

To isolate exogenous shifts in the marriage market equilibrium, I first show that positive rainfall shocks in a girl's pre-marital adolescence increase marriage inequality in polygamous marriage markets and reduce it in monogamous ones. Across several specifications, which control for demographics, spatial characteristics, current and lagged rainfall, and ethnicity and region fixed effects, a 1 SD increase in annual rainfall from its long-run mean in the adolescent period of the average woman differentially increases the Gini coefficient of brides in polygamous villages by 9.9-12.8 Gini-points, and reduces the marriage rate for men by 9.4-14.2 percentage points. This differential inequality response is driven by a reallocation of marriage market share toward older men in polygamous villages, and younger men in monogamous ones.

These aggregate market effects are driven by marriage decisions at the household level. Consistent with Corno et al. (2017), positive pre-marital income shocks increase age of marriage and reduce child marriage. Women who experience good shocks as adolescents also exhibit lower annual hazard into marriage over their peak marriageable years, an effect driven entirely by polygamous communities. To probe the robustness of these results, I conduct placebo tests which reveal that pre-adolescent rainfall shocks are not significantly related to marriage decisions. Taken together, the results imply that plausibly exogenous fluctuations in historical local weather conditions change families' marriage choices and equilibrium outcomes in the marriage market.

I then use these results as the first stage in an instrumental variable strategy to estimate the effect of marriage inequality on Boko Haram activity. Marriage inequality increases from an income shock only in polygamous markets, allowing me to use the interaction between village-level polygamy status and pre-marital rainfall shocks as an instrument for inequality. In the reduced form, I estimate that a 1 standard deviation increase in rainfall during the adolescent period of the average woman leads to a differential increase of 10-15 Boko Haram-related deaths annually in polygamous villages. In the 2SLS models, a one-point increase in the Gini coefficient of brides increases Boko Haram fatalities by an 0.7-1.6 annual deaths. The first stage, reduced form, and 2SLS estimates are all conditional on current and lagged rainfall, ruling out direct effects of current climate conditions on marriage markets and conflict.

To explain the results, I develop a simple model of occupational choice and marriage in which men choose to enter the civilian marriage market or an join an extremist group. The key prediction of this model is that good economic conditions increase violence under polygamy. Good income shocks in a girl's adolescence improve her family's outside option, increasing the reservation bride price and delaying marriage. Rising bride prices push poorer men out of the market and women toward richer men, increasing inequality in the distribution of brides. Positive income shocks reduce the marginal value of today's dollar via diminishing marginal utility, thereby increasing the sensitivity of reservation prices to the size of tomorrow's payoff. Since polygamy increases the option value of waiting – as women have the chance of matching with wealthier married men – it therefore increases both the initial reservation bride price and its derivative with respect to outside income. Violence increases as young men at the margin of indifference choose insurgency. The marginal reduction in marriage value is greater in polygamous areas, so these areas experience more severe outbreaks of violence in response to good pre-marital economic conditions.

In an extension to the model in which families have preferences over both the age and bride price of their daughter's matches, I show that positive shocks can actually reduce marriage inequality under monogamy, consistent with the empirical results. As reservation utility rises with income, women in polygamous areas require higher bride prices that can only be afforded by married older men, while those in monogamous areas prefer to adjust along the age margin since older unmarried men in monogamous areas are negatively selected. Under monogamy, a positive female income shock is accompanied by reallocations of marriage market share from older to younger men, while the opposite occurs in under polygamy. This pattern is observed in the data.

I then turn to explore the mechanisms underlying the main results. Using data on household expenditures, I find that average wedding expenditures – a proxy for bride prices – also increase in response to the adolescent rainfall shock. A 1 SD deviation from the long-run mean of rainfall in the adolescent period of the average woman increases marriage expenditures by 70-90% in polygamous markets, but has no impact in monogamous ones. Therefore, the reduction in marriage rates and increase in market concentration observed in the first stage regressions is likely driven by reservation prices, rather than by other marriage market primitives such as preferences, institutions, sex ratios, or male income. Consistent with this, I observe that women who experience positive shocks as adolescents marry richer and more polygamous men, which suggests a stronger market position and is consistent with the observed market-level increase in inequality. Again, this effect is concentrated in polygamous areas. Finally, I find that abductions and violence against women by Boko Haram increase in polygamous areas in response to adolescent rainfall shocks. This suggests that Boko Haram responds strategically to marriage market imbalances by increasing abduction of adolescent girls to meet demand for wives among new recruits.

The results are robust to a wide variety of falsification tests, sensitivity analyses, and controls. While the problems in using rainfall as an instrument are well known (Sarsons 2015), my IV uses a female-specific cohort-weighted average of shocks that should be uncorrelated with current income, conditional on current rainfall. To support the exclusion restriction, I conduct balance tests which demonstrate that numerous covariates are uncorrelated with the interacted IV. I test the robustness of all models to controls for wealth inequality, ethnic and religious group characteristics and their interactions with rainfall shocks, and current and lagged rainfall interacted with polygamy to rule out bias induced by serial correlation in rainfall. Several placebo tests confirm that pre-adolescent rainfall shocks – which have no marriage market effects – have small or null conflict effects. In Appendix D I additionally probe robustness to different subsamples, definitions of marriage inequality and polygamy status, and conflict data sources, demand-side changes, account for male and female migration, and conduct Bayesian IV diagnostic exercises.

This paper contributes to three literatures in economics. The first is the empirical literature on economic shocks and civil conflict. It is well known that violence responds to economic incentives. Economic shocks that improve welfare typically raise opportunity costs of fighting, reducing violence (Humphreys and Weinstein 2008, Bazzi and Blattman 2014, Hodler and Raschky 2014, McGuirk and Burke 2017, Gehring et al. 2018, Blattman and Annan 2016). On the other hand, some positive shocks – such as rising commodity prices – may both improve welfare and increase the value of capturing the state, generating a countervailing "rapacity" effect that increases violence (Dube and Vargas 2013,Nunn and Qian 2014, McGuirk and Burke 2017, Angrist and Kugler 2008). The focus of much of the empirical literature on the economics of conflict has been adjudicating these two forces. This paper applies a similar opportunity cost-based argument. However, I show that comparative statics of the opportunity cost mechanism can be reversed when the marriage market is accounted for: positive income shocks can exacerbate conflict by increasing marriage inequality when marriage markets exhibit polygamy. This is a novel finding, which points to the importance of traditional cultural practices in determining the effects of economic shocks.

The findings also speak to the literature on marriage markets and crime, where researchers have typically studied the role sex ratio imbalances (Edlund et al. 2013, Cameron et al. 2017). I expand this literature by providing more detailed mechanisms underlying the link between marriage market fluctuations and violence, identifying how institutions and outside options influence inequality, the value of marriage, and the opportunity costs of violence. I also consider forms of violence beyond criminality and intimate partner violence, suggesting that marriage market causes may underlie a wide range of violence, from crime to religious extremism.

Lastly, the results add to the emerging literature on the welfare consequences of traditional marriage practices. Since Becker (1981) and Grossbard (1980), economists have studied polygamy and bride price as equilibrium outcomes, but only recently has an emerging empirical literature begun to investigate the implications of these practices for outcomes such as child mortality (Arthi and Fenske 2018), child marriage (Corno et al. 2017 and Corno and Voena 2016), investment in girls' education (Ashraf et al. 2016), and marital quality (Lowes and Nunn 2016). This is the first paper to rigorously identify the link between polygamy and civil conflict and demonstrate the underlying mechanisms, existing evidence being found only in cross-country regressions (Kanazawa 2009). More broadly, the results speak to importance of culture and traditional practices in shaping economic outcomes, conflict, and development, as in recent work by Moscona et al. (2018).

The paper proceeds as follows. Section 2 provides background information on traditional marriage practices in Nigeria and the history of the Boko Haram insurgency. Section 3 lays out a model of the marriage market and occupational choice to motivate the empirical analysis. Section 4 describes the data and presents summary statistics. Section 5 details the empirical strategy and provides preliminary tests to validate the instrument. Section 6 provides the primary results and robustness tests, while section 7 provides results illustrating key mechanisms. Section 8 concludes.

2 Background

2.1 Marriage customs in Nigeria

Marriage markets in Nigeria are characterized by two key institutions: bride price and polygamy. Bride price is a marriage practice that dictates payments from the groom to the family of the bride at the time of marriage, although allocation of property rights over this transfer between bride and her parents varies across societies (Anderson 2007). Historically, nearly all of the ethnic groups in Nigeria have practiced both polygamy and bride price. Data from Murdock (1967), which codes an extensive set of traditional practices across 1265 pre-modern societies, provide insights about traditional marriage practices across 95 Nigerian ethnic groups observed at varying points between 1870-1950. Every ethnic group in this sample practices some form bride price, while 86% practice direct monetary transfers to the bride's family.[4] At the same time, only two groups in the Nigeria sample, the Chamba and the Jibu, practice strictly monogamous marriage.

[4] The remainder also practice bride price in the form of either token or in-kind transfers.

Despite the lack of group-level differences in historical marriage institutions, there is variation in polygamy rates across Nigeria. As Fenkse (2015) demonstrates using data from Demographic and Health Surveys (DHS), there is substantial within-country variation in rates of polygamous marriage across Africa. In my sample, there is variation in polygamous practice both at the household-level within communities, and across villages and regions. To capture variation in polygamy as a long-standing norm rather than the polygamy rate itself, which is an equilibrium outcome of the marriage market, throughout the paper I define polygamy as a dummy that indicates whether a village contains any polygamous man. By this measure, polygamy is practiced in 61% of villages included in the 2008 and 2013 Nigeria DHS rounds. While this variable may be correlated with cultural differences across tribes, in villages without any observed polygamy it is likely that all households are far from the margin at which variation in unobservables or bride prices would push them into or out of polygamy. In Appendix Table D2, I show that all results are robust to using different thresholds to define cluster-level polygamy status, as well as using equivalent definitions based on the larger woman-level sample.

There is qualitative evidence that high bride prices and polygamy in Nigeria have led to concentration of brides in the hands of richer, older men (Hudson and Matfess 2017). Two facts are consistent with this story: *i*) large gender gaps in age of marriage suggests that men must accumulate significant wealth in order to afford marriage, and *ii*) a strong within-village correlation between wealth, age and brides,[5] which suggests that polygamy and high bride price concentrate wives among older and richer men. Ethnographically, Cohen (1961) finds that polygamy among the Kanuri in Northeast Nigeria[6] engenders a status hierarchy among men based on the number of wives. A concentration of young men at the bottom of this hierarchy may brew social instability.

2.2 Boko Haram

Boko Haram is an Islamist terror group espousing Salafi jihad and the violent imposition of an Islamic state in Northern Nigeria. It was founded in 2002 by the itinerant preacher Mohammed Yusuf in Maiduguri, the capital of Borno State. Though the group has always espoused violent jihad, it did not begin active operations until 2008. In 2009, Boko Haram announced their presence in a wave of violence that swept through Northern Nigeria, killing over 1000 civilians.

Since 2009, Boko Haram has been engaged in near-constant conflict with the Nigerian military and has preyed on local civilian populations. The insurgency has resulted in a humanitarian emer-

[5] These results are presented in Table C2 and discussed in detail in Appendix C.
[6] The Kanuri homeland makes up the center of the Boko Haram insurgency.

gency, including the death of more than 20,000 people and the displacement of roughly 1.7 million others.[7] At its height, the group controlled significant territory in Northern Nigeria and boasted an estimated fighting force of roughly 15,000.[8] The group is known for its extreme brutality, its use of mass abductions, and frequent use of women and children as suicide bombers.

Interviews with militants suggest that recruitment is driven by poverty and unemployment in Northern Nigeria (Onuoha, 2014). However, marriage may be a potent force in funneling recruits to Boko Haram. Several observations suggest that this channel is plausible. Boko Haram is distinguished from other jihadist groups by its use of mass abductions of schoolgirls,[9] which suggests that controlling large numbers of adolescent girls is strategically important to the group. Boko Haram has explicitly used offers of marriage to attract young men, and bride price emerges as a key concern among young members (Cold-Ravnkilde and Plambech 2015, Hudson and Matfess 2017). Unlike other terrorist groups in which captive women are an ancillary membership benefit, qualitative evidence shows that young men in Northern Nigeria do join Boko Haram for marriage: women are abducted for this purpose, trained as wives, and men are rewarded for their service with affordable, recognized marriages (Hudson and Matfess 2017).

3 Model

In this section, I present a model that derives marriage and occupational decisions as a function of exogenous economic conditions and marriage market institutions. In the model, men make bride price offers based on their position in the income distribution, and a girl's parents accept or reject these offers based on outside options. The resulting marriage market equilibrium determines the expected value of civilian life, which is traded off by young men against the value of joining insurgents. The key exogenous parameters in the model are the marriage regime – monogamy or polygamy – the distribution of male types, and the girl's contribution to household income. The endogenous variables are the equilibrium bride price, the marriage rate of different male types, and the probability of joining Boko Haram. The model delivers several testable predictions: *i*) that polygamous marriage markets have greater marriage inequality, *ii*) that marriage inequality increases in response to positive female income shocks, more so under polygamy, and *iii*) that positive female income shocks increase Boko Haram activity, more so under polygamy.

[7] https://www.unocha.org/story/five-things-know-about-crisis-nigeria

[8] https://www.amnesty.org/en/latest/news/2015/01/boko-haram-glance/

[9] See for example, the Chibok girls http://www.bbc.com/news/world-africa-32299943, or the more recently in Dapchi, Yobe State https://guardian.ng/news/bring-back-our-girls-blames-government-failures-for-dapchi-kidnap/

3.1 Set up

Consider a very simple marriage market with polygamy, a reformulation of the classic Becker (1974) model with dynamic considerations, as in Corno et al. (2017). There are four types of men who differ along two dimensions: they are young or old $\{Y, O\}$ and they are rich or poor $\{R, P\}$. In a monogamous marriage market, only the the young participate, while in a polygamous market, the rich old also participate, capturing the empirical association between income and wives. There are N^F women and N^j men of each type. Throughout, I assume equal population shares between the types. Men have values for marriage as follows: $v^{RO} > v^{RY} > v^{PY}$. Each period every woman meets a randomly drawn man with probability λ, and meets nobody with probability $1 - \lambda$. If a meeting happens, the man of type j makes an offer w^j and gets surplus $v^j - w^j$.

Following the qualitative literature (Grossbard 2015, Erulkar and Bello 2007), I assume a girl's parents make decisions about her marriage offers. They have log utility. Girls generate income y^f for their family each period if they remain unmarried. Parents discount the future by a factor β. Given an offer w, they solve the discrete problem

$$\max\{\log(w), \log(y^f) + \beta V_r^{wait}\}$$

Where V^{wait} is the value of staying in the market until the next period. r is the regime, either monogamous m or polygamous p. Before the marriage market, young men must decide whether to enter civilian life or join a rebel group, which is absorbing. The rebels offer a marriage package worth v^{reb} to all young men. If the man enters the rebel world, he gains the offer of marriage v^{reb} and incurs disutility ξ_i of rebel activity, drawn from distribution Ξ, which is independent of his type. The utility from taking the rebel offer is $v^{reb} - \xi_i$.

3.2 Marriage market

First note that parents will follow a cutoff strategy, accepting w such that $\log(w) > \log(y^f) + \beta V_r^{wait}$ and declining otherwise. The reservation offer $\bar{w}$ will therefore satisfy

$$\bar{w} = y^f \exp(\beta V_r^{wait})$$

I assume that the value of waiting in a polygamous market is $V_p^{wait} > V_m^{wait}$. This is because additional high-value men in the market both increase the likelihood of meeting a match next period, and since these men are rich type, increase the expected offer. Both factors increase the expected value of next-period matches. Together, this yields the first result: that polygamy creates greater reservation bride prices. The market under polygamy and monogamy is illustrated in Figure 1. Demand forms a step function, with the steps indicated by the values of the different types of men. In the monogamous market (Panel A), only RY and PY types enter the market, while in the polygamous market (Panel A), demand begins at the higher v^{RO} due to older, richer men re-entering the marriage market. Supply is perfectly elastic at $\bar{w}$ until N^F, after which all women have entered the market and it becomes inelastic. The reservation price under polygamy is higher than under monogamy because of the difference in continuation values.

Figure 1

To begin, I compare the properties of the initial equilibria under polygamy and monogamy. First, I assume $N^Y < N^F < N^Y + N^{RO}$ because child marriage increases the supply of women. The first result is that in polygamous equilibrium E_p, prices are higher than in the monogamous equilibrium at E_m. This is driven by two forces: *i*) the reservation bride price is higher under polygamy, and *ii*) the entry of RO-type men shifts out demand and drives up the equilibrium price from $\bar{w}_p$ to v^{PY}. The second result is that marriage inequality is higher under polygamy. At E_m, everyone is married to exactly one woman, and there is perfect equality. But at E_p, some PY are unmarried, while all RO have 2 wives. Finally, note that both of these predictions depend in magnitude, but not direction, on elasticity of supply. If supply is perfectly elastic, then the polygamous equilibrium is $\tilde{E}_p$, while the monogamous one is unchanged. Equilibrium polygamous bride prices fall to $\bar{w}_p$, still higher than under monogamy. Wives are now affordable to all men, so inequality falls. Still, inequality persists under polygamy because RO men have two wives.

The key result is that positive female income shocks increase the reservation price and leading to greater marriage inequality only under polygamy. To see this, note that $\frac{\partial \bar{w}}{\partial y^f} = \exp(\beta V_r^{wait}) = \frac{\bar{w}}{y^f}$. The derivative is proportional to the initial reservation wage, and so greater under polygamy. In Figure 1, this corresponds to an upward shift in the equilibria to E'_p and E'_m, respectively. Under monogamy, the shift in reservation prices is small enough that all men remain married. However, given that $\frac{\partial \bar{w}}{\partial y^f}$ is bigger under polygamy, an equivalent Δy^f increases prices enough to push all of the PY types out of the market. Thus inequality is unchanged in m but rises in p. The relative magnitudes of bride price effects are ambiguous: if supply is perfectly elastic, or initial $\bar{w}_p$ is very large relative to $\bar{w}_m$, then price increases will be greater under polygamy. However, as drawn, the price increase is lower under polygamy, where it shifts only from v^{PY} to $\bar{w}'_p$.

3.3 Occupational choice

Now consider the decision of a young man who is considering entering the civilian or militant market under regime r before observing his type. His expected value of civilian life will be

$$\begin{aligned} V_r^{civ}(w_r) = \lambda \frac{N^{RY}}{N^Y} & \left[(v^{RY} - w_r) 1(v^{RY} \geq w_r) + 1(v^{RY} < w_r) \beta V_r^{search,RY} \right] \\ & + \lambda \frac{N^{PY}}{N^Y} \left[(v^{PY} - w_r) 1(v^{PY} \geq w_r) + 1(v^{PY} < w_r) \beta V_r^{search,PY} \right] \\ & + (1-\lambda) \beta V_r^{search} \end{aligned}$$

With probability λ, he meets a match and draws his type. The expectation is taken over the two possible states, rich or poor, comprising the first two terms. In either case, if the equilibrium bride price w_r is below his value v^j, he obtains marriage value $v^j - w_r$; if above, he continues to search conditional on type. With probability $1 - \lambda$, he obtains the continuation value. Assuming a stationary equilibrium, as shown in Appendix E the value function can be re-written

$$V_r^{civ}(w_r) = \frac{\lambda}{1-(1-\lambda)\beta} \left[\frac{N^{RY}}{N^Y} (v^{RY} - w_r) 1(v^{RY} \geq w_r) + \frac{N^{PY}}{N^Y} (v^{PY} - w_r) 1(v^{PY} \geq w_r) \right]$$

In contrast, his value from joining a militia will be $v^{reb} - \xi_i$. The joining rate ρ_r will be

$$\rho_r = \Xi \left(v^{reb} - V_r^{civ}(w_r) \right)$$

The joining rate – and Boko Haram violence – is increasing in y^f for both r

$$\begin{aligned} \frac{\partial \rho_r}{\partial y^f} &= \frac{\partial \rho_r}{\partial V_r^{civ}} \frac{\partial V_r^{civ}}{\partial w_r} \frac{\partial w_r}{\partial y^f} \\ &= -\Xi' \left(v^{reb} - V_r^{civ}(w_r) \right) \frac{\partial V_r^{civ}(w_r)}{\partial w_r} \frac{\partial w_r}{\partial y^f} > 0 \end{aligned}$$

Since $\Xi' > 0, \frac{\partial V_r^{civ}(w_r)}{\partial w_r} < 0, \frac{\partial w_r}{\partial y^f} > 0$. The final prediction is that $\frac{\partial \rho_p}{\partial y^f} > \frac{\partial \rho_m}{\partial y^f}$. To see this, consider the simplest case of perfectly elastic supply, $w = \bar{w}$. We have already established that the term $\frac{\partial \bar{w}}{\partial y^f}$ is larger under polygamy. Assuming equal group sizes, $\frac{\partial V_p^{civ}(\bar{w}_p)}{\partial \bar{w}_p} = \frac{\partial V_m^{civ}(\bar{w}_m)}{\partial \bar{w}_m}$. Finally, $\Xi' \left(v^{reb} - V_p^{civ}(\bar{w}_p) \right) > \Xi' \left(v^{reb} - V_m^{civ}(\bar{w}_m) \right)$ as long as ρ_r is sufficiently low for both r. This depends on the distribution for ξ. If Ξ is symmetric, then using $V_p^{civ}(\bar{w}_p) < V_m^{civ}(\bar{w}_m)$, we have that $v^{reb} - V_p^{civ}(\bar{w}_p) > v^{reb} - V_m^{civ}(\bar{w}_m)$. As long as $\rho_p \leq 0.5$,[10] the pdf is increasing and the inequality

[10] This seems reasonable given that the number of Boko Haram fighters is small relative to the population.

holds. An increase in right-skew will lower the threshold for ρ. A positive female income shock makes men more likely to join Boko Haram by raising the reservation bride price and reducing the expected value of civilian life under polygamy.

The conclusions of the model rest on a number of assumptions. The assumption that older and richer men have more willingness-to-pay is key to generating greater marriage inequality in polygamous places. Appendix C.2 demonstrates that individual wealth and age are both positively correlated with the number of brides. Parents are assumed to have diminishing marginal utility; under linearity, the sensitivity of the reservation price to outside income no longer depends on the regime. More broadly, the shape of the supply curve matters. Under perfectly elastic supply, bride price increases from a female income shock are always greater under polygamy. With fixed supply, this prediction depends both on the sex ratio and the gaps in the male utilities. Another key assumption is regime-dependent continuation values, which drives the dynamics of reservation bride prices; greater polygamous demand should naturally increase both next-period prices and match probability for each period. Skewness in the disutility pdf determines the maximal joining rate that is sufficient to make the final Boko Haram result. If the pdf is monotonically decreasing, there can be no sufficiency condition on the distribution, ruling out e.g. power laws.

4 Data and summary statistics

4.1 Data description

To test the key hypotheses about extremist violence, marriage, and female income, I require community-level data on exposure to Boko Haram, marriage market statistics such as the polygamy rate, mean number of wives, and marriage inequality, individual-level demographic data, and a full rainfall history for each location to construct rainfall shocks.

Marriage data: Data on the marriage market comes from the Nigeria Demographic and Health Surveys (DHS). DHS is a nationally representative survey conducted every five years and includes a household questionnaire, as well as individual interviews with a sample of men aged 15-59 and women aged 15-49. In addition to a range of data on marriage, reproduction, education, and health, DHS also provides GPS locations at the cluster-level–the primary sampling unit–which allows me to match clusters to both violent events and rainfall histories.[11] I use repeated DHS cross-sections for

[11] Cluster GPS coordinates are displaced by 10-15 km to preserve the anonymity of respondents. This displacement is random, introducing classical measurement error in both dependent and independent variables.

2008 and 2013, discarding earlier data that predates the existence of Boko Haram.[12]

Using responses on the number of wives a man has, I measure marriage-market inequality by calculating the Gini coefficient of wives for men aged 15-59. All calculations include zero values for unmarried men. I also take the following variables from the men's, women's, and household surveys: number of wives, a polygamy dummy, self-reported ethnicity, religion, age of marriage, year of birth, and household wealth measured on a continuous index scale ranging from -3 to 3, though I re-center this variable to begin at zero. Combining the 2008 and 2013 data, I obtain a final sample of 32,497 men, 72,009 women, and 1,768 clusters for which all relevant data is available[13]. I use clusters as the unit of analysis for all market-level regressions, implicitly defining a marriage market as a cluster.[14] However, I test the robustness of this assumption in Appendix D.9. In all market-level regressions, I collapse variables measured at the individual-level to their cluster-level means, and apply cluster-level sampling weights.

Figure A1 maps the geographic distribution of the male polygamy rate across DHS clusters in the 2008 (Panel A) and 2013 (Panel B) DHS rounds. In both rounds, polygamy appears to be broadly distributed across the country, obviating concerns that the correlation between local ethnic and religious heterogeneity might be contaminating our results. If anything, polygamy appears relatively high in Nigeria's south and Northwest, but relatively low in the Northeast, the epicenter of the Boko Haram rebellion.

Violence data: Data on Boko Haram incidents comes from the Armed Conflict Location Event Dataset (ACLED), a standard event dataset used throughout the political economy and conflict literature, which identifies events using local and international newspapers (see Raleigh et al. 2010 for a description). I identify 2,016 Boko Haram-involved events, roughly 20% of Nigeria's total number of violent events. I match these events to DHS clusters by assigning to each cluster all of the events that fall within a 20 kilometer radius. I then consider only the events that occur between DHS rounds to capture the cross-section of conflict in a given cluster at a given time. That is, for 2008 I retain events occurring from 2008-2012, and for 2013 I retain events occurring from 2013-2016.[15] I measure conflict as the mean annual number of events or fatalities. Recent methodological papers have

[12] However, for testing hypotheses that do not involve the Boko Haram data, I also include DHS data from 2003 to maximize power.

[13] Women are oversampled in the DHS to estimate various maternal and child health indicators with greater precision. This sample size differential does not affect measures of marriage inequality, since I use only the male data to construct these variables

[14] While marriage markets in remote rural Nigeria are likely to be spatially isolated, this may not be the case in more urbanized areas. However, to the extent that shocks affect marriage outcomes in clusters in which they did not occur via spatial integration, this should bias the results toward zero.

[15] Results are generally similar if limited to 2 and 3 year periods. Results available upon request.

raised concerns over underreporting and other reliability issues in media-based event data (Weidmann 2015,Eck 2012). However, under the plausible assumption that any reporting biases in the ACLED data collection process are uncorrelated with cluster-level mean adolescent rainfall shocks, these measurement problems should not affect the results.

Figure A2 maps the geographic distribution of mean annual Boko Haram fatalities across DHS clusters for the two key periods in our data: 2008-2012 and 2013-2016. In the first period (Panel A), most of the Boko Haram deaths are clustered in the Northeastern corner of Nigeria, primarily in Borno state. Maiduguri, the home of the group, has the highest density of Boko Haram fatalities. While attacks appear to be more sparsely distributed across the rest of Northern Nigeria, only Kano, the largest city and commercial capital of Northern Nigeria, has a substantial density. Interestingly, these patterns change when considering Panel B (2013-2016). In the later period of Boko Haram's reign of terror, attacks appear to be far less concentrated in Borno state, and more evenly distributed across all of Northern Nigeria, even reaching Sokoto in the far northwest, which had been spared from the group's earlier activity. This suggests an expanding reach and perhaps increased recruitment activity in this period, even as the group's territorial control has been pushed back to its heartland on the Cameroonian border and the Lake Chad basin (The Guardian 2016).

Rainfall data: Following the standard in the climate economics literature, I take rainfall from the University of Delaware (UDEL) gridded precipitation dataset (Dell et al., 2014). The data are gridded at a spatial resolution of 0.5×0.5 degrees (roughly 50 kilometers at the equator), and are interpolated from weather station readings. For each grid point, the data contain monthly precipitation levels from 1900 to 2014, as well as the coordinates of the centroid of the grid. I match the grids to DHS clusters by finding the nearest grid centroid to each cluster, attaching the rainfall history to the DHS cluster. I obtain 287 rainfall grids matched to any DHS cluster.

Given the size of the rainfall grids, several DHS clusters can be assigned to the same grid-cell. This results in spatial autocorrelation within cells–each cluster within a grid contains the exact same rainfall history.[16] I address this issue by clustering standard errors at the grid-cell level.[17] I use rainfall data for two purposes: firstly, to construct the deviation of rainfall from its long-run mean experienced during a woman's pre-marital adolescent years, which I take to be ages 12-16, following previous literature (Fenkse 2013 Corno et al. 2017). This forms the key input into my cluster-level adolescent rainfall shock. I also use the rainfall data to control directly for current and lagged rainfall

[16] However, this does not imply perfect within-grid correlation in the shocks I use for identification, since the construction of these shocks also exploits cross-sectional variation in female cohort composition across clusters.

[17] Note, however, that this does not address broader issues of spatial autocorrelation induced by the use of interpolation methods with sparse weather station data

shocks in the survey year in order to address endogeneity concerns.

Additional data: In several specifications, I include geographic fixed effects that control for unobserved heterogeneity in the state of residence and ethnic homeland. The former accounts for state-specific differences in policies, governance, or level of development; the latter controls for tribe-specific cultural and geographic factors that may be correlated with Boko Haram activity and marriage market variables. I obtain the boundaries of ethnic homelands from the Murdock (1967) *Ethnographic Atlas*, a comprehensive map of the homelands of ethnic groups across sub-Saharan Africa which has been used extensively in empirical research African economic history, digitized by Nunn (2008). I assign each DHS cluster to the ethnic homeland within which it falls. I view including these homeland fixed effects as complementary to controlling directly for ethnic shares with the self-reported DHS data, since they are less likely to suffer from reporting bias, but also too blunt to capture remaining cluster-level heterogeneity in ethnic composition within an ethnic homeland area. Finally, I obtain controls for fixed cluster characteristics such as malaria exposure, slope, mean monthly temperature in survey year, modeled population density, and others from the DHS geospatial data files, which include an extensive set of spatial variables.

4.2 Descriptive statistics

Summary statistics for the data at cluster and individual-level are presented in Table 1. I estimate all summary statistics by DHS year and marital regime r, as well as include full-sample totals for each year. In the top panel, I estimate means and standard deviations of the cluster-level variables used in the bulk of the analysis: marriage inequality, wealth inequality, Boko Haram activity, population density, and an indicator of whether the cluster contains any polygamous men. A few points are salient. Firstly, marriage inequality is high across Nigeria, and rising, with an average within-cluster bride Gini of 0.5 in 2008 and 0.54 in 2013. In contrast, average within-village wealth inequality is low and falling, from 0.17 to 0.1. Polygamy is a common institution with around 60% of villages containing at least one polygamous man, a proportion that falls slightly over time. Boko Haram activity is widespread and rising, affecting 18.5% of clusters in 2008 and 25.9% in 2013. Average annual fatalities per cluster also rise over time, from 3.29 to 11.95, reflecting an increase both in the level of activity and deadliness over this period. Still, Boko Haram-related kidnappings remain roughly constant over time, affecting 3-4% of villages.

Table 1

Across cluster-level observables, monogamous and polygamous areas are generally quite similar. In both rounds, they have comparable levels of marriage and wealth inequality, as well as Boko Haram presence. Only when it comes to mean annual Boko Haram deaths – the main outcome of interest – do they exhibit large differences, present in both rounds. Monogamous villages are also substantially more dense, indicating that they are likely less remote. This is consistent with higher levels of average wealth in monogamous villages in both rounds.

The second and third panels summarize marriage and demographic data for the 32,497 men and 72,009 women in the sample.[18] There is a large gap in average age of marriage, with men marrying around age 25 on average, while the average woman marries just before age 18, below the minimum lawful age of marriage nationally. Among men aged 15-59, 45-50% are unmarried, 40-45% are monogamously married, while the rest (roughly 10-15%) are polygamous. Therefore, polygamists comprise roughly 20% of married men. While these are relatively small percentages of men, they may be able to generate effects on marriage inequality and conflict for two reasons: Firstly, since Boko Haram can do substantial damage with even a small number of recruits, if any of the men displaced from the market by polygamy are marginal to joining, it can have a outsize effect on conflict. Secondly, as the model makes clear, it is the existence of a polygamy norm itself, rather than the equilibrium level, that is sufficient to generate the dynamics of the model. Among women, there are higher rates of polygamy–roughly 33% of married women aged 15-49 are in polygamous unions, and these women have an average of 1.4 co-wives. Polygamy rates for men and women do not vary substantially over time, suggesting a stable equilibrium.

Demographics are highly correlated with marital regime in this simple comparison of means. Polygamous villages have larger shares of Muslim and Hausa inhabitants, on average, suggesting that I should test robustness of the main results to omitted ethnoreligious characteristics (as in Table B5). As expected, age of marriage also appears to vary with polygamy status; it is roughly three years lower for men and nearly 3.5 years lower for women in polygamous areas. This is consistent with both polygamous areas having stronger traditional mores or greater marriage market surplus incentivizing earlier entry.

The distributions of key cluster-level variables – marriage inequality, adolescent rainfall, polygamy, and average number of wives – are seen in Figure A3. In Panel B, it is worth noting that most of the mean rainfall deviations fall below zero, suggesting that, on average, bad adolescent rainfall realizations relative to the long-run mean are more common than good ones. Also note that the distribution

[18] Sample sizes differ from those mentioned in Section 4 because Table 1 uses the final estimation sample, which drops observations with missing data.

of village-level shocks ranges from -1 to 0.5; a 1-unit change in the key independent variable corresponds to moving from the first to the 90th percentile of the distribution, important to bear in mind when interpreting the coefficients of the reduced form and first-stage regressions. Marriage inequality is fairly symmetrically distributed, while both polygamy rate and mean number of wives have a point mass at zero and one, respectively, and long right tails.

As predicted by the model, the equilibrium relationship between polygamy and marriage inequality is strongly positive. In Figure A4 I plot the binned partial correlation between the village-level marriage Gini and male polygamy rate, controlling for DHS round fixed effects, ethnicity fixed effects, and a large set of control variables. This is consistent with the insight that polygamy pushes up reservation prices via increased option value, cutting low type men out of the market; at the same time, it allows high-type men to have multiple wives. Both of these forces exert upward pressure on inequality. Despite this strong correlation, the endogeneity of the polygamy rate renders it unsuitable as an instrument for marriage inequality.

Is marriage inequality correlated with Boko Haram activity? In Figure 2 I plot both the unconditional and conditional correlations in a binned scatterplot, with thirty bins at percentiles of the cluster-level distribution of bride Gini coefficients, for both polygamous and monogamous clusters.[19] Panels A and C give the unconditional correlation, which is positive in polygamous clusters and flat in monogamous ones. The conditional relationship, which controls for the full set of controls, tribe fixed effects, and DHS round fixed dummies, is given in Panels B and D. In both sub-samples there is a positive relationship, stronger than the unconditional specification. When I use the full sample, the unconditional relationship has a slope of 0.25, significant only at the 10% level, while the conditional slope estimate is 0.43, significant at the 5% level. While the positive association is somewhat robust, I defer a discussion of causality to Section 6.

Figure 2

5 Empirical strategy

5.1 Main specification

The model predicts that higher female pre-marital income y^f reduces the marriage rate of low-type men, and this effect is greater under polygamy. This in turn reduces the expected value of civilian life V^{civ} and increases ρ, the recruitment rate. The empirical implication is that the interaction

[19] The slopes are estimated on the underlying cluster-level data.

between income shocks and polygamy should correlate positively with marriage inequality and Boko Haram activity. The model lends itself naturally to an IV interpretation; the equation for the marriage market value forms a first stage for the occupational choice equation. The interaction of polygamy status with female income can serve as an instrument for marriage market conditions to estimate the effect of the marriage market on Boko Haram activity. The estimating equation corresponding to the occupational choice equation, for cluster j in survey round r is

$$B_{jr} = \alpha_0 + \varphi G_{jr} + \delta_r + \gamma_s + \xi_g + X'_{jr}\beta + u_{jr}$$

Where B_{jr} are Boko Haram-related deaths in cluster j in the years between round r and $r+1$. G_{jr} is the Gini coefficient of brides in cluster j among all sampled men aged 15-59. γ_s are state fixed effects and ξ_g are ethnic homeland fixed effects. The coefficient of interest is φ, which represents the marginal change in Boko Haram-related fatalities corresponding to a one point increase in the bride Gini. The parameter φ is estimated via 2SLS with the following first stage:

$$G_{jr} = \alpha_1 + \pi_1 p_{jr} s_{jr} + \pi_2 p_{jr} + \pi_3 s_{jr} + \delta_r + \gamma_s + \xi_g + X'_{jr}\beta + \varepsilon_{jr}$$

Where p_{jr} is a dummy for polygamy status, and s_{jr} is the average deviation of rainfall from it's long-run mean during the adolescent period of women in cluster j. The coefficient of interest in the first-stage is the interaction term, π_1. This captures the differential effect of the rainfall shock on marriage inequality in polygamous societies, relative to monogamous ones. This coefficient serves both as a test of the theoretical model, which predicts $\pi_1 > 0$, and as identifying variation for the IV models. Throughout, I measure p_{jr} as a dummy variable that equals one when any household in village j is polygamous. Note that both p_{jr} and s_{jr} are included in the second stage, along with all of the fixed effects. Only their interaction is an excluded instrument.

All regressions include a vector of cluster-specific controls X_{jr}. Unless otherwise specified I typically include DHS round fixed effects, annual rainfall deviation and average monthly temperature in the year of the survey, lagged rainfall deviation, cluster gradient, population density in 2005, average age of men and women, wealth inequality, distance to borders, share Muslim, and share Hausa, the dominant ethnic group in Northern Nigeria. Standard errors are clustered at the grid-cell level to account for spatial correlation in shocks across clusters within grids. This should provide somewhat conservative estimates, since each cluster is matched to a larger rainfall grid.

5.2 Measurement assumptions

A few assumptions are required to implement this strategy. Firstly, the joining rate ρ is unobserved. However, if we assume that the level of violence Boko Haram perpetrates in a given locality is an increasing function of the number of recruits, then we can use observable conflict intensity as our key outcome measure. In Appendix D.10 I interrogate the validity of this assumption in greater detail. We also do not observe V^{civ}, the value of the civilian marriage market, which links the exogenous features of the marriage market to Boko Haram activity. As a proxy, I use the level of inequality in the distribution of brides and the male unmarried rate, both easily observed indicators of marriage market tightness.

Following the long literature in development economics begun by Miguel et al. (2004), I use rainfall to measure y^f, the economic conditions in which families form their reservation bride prices. In particular, I construct cohort-and-cluster specific rainfall shocks that measure the deviation between the average annual rainfall over a woman's pre-marital adolescence, defined as ages 12-16, and the cluster's long-run average. This is the relevant deviation for the parents' reservation price, as it represents the economic environment as the girl enters the marriage market. Let μ_j be the long-run average rainfall in cluster j, and σ_j be its standard deviation. r_{jt} is the rainfall experienced by cluster j in year t. For woman i in cluster j born in year τ_i (alternatively, belonging to cohort c), the individual rainfall shock will be

$$s_{ij} = \frac{\frac{1}{T}\sum_{t=\tau_i+12}^{\tau_i+16} r_{jt} - \mu_j}{\sigma_j}$$

Therefore, the cluster-level average, s_j, can be written as

$$s_j = \frac{1}{N_j}\sum_{i=1}^{N_j} s_{ij} = \sum_{c=1}^{C} s_{cj}\omega_{cj}$$

The cluster-specific shock s_j represents the mean of individual shocks s_{ij}, where each s_{ij} is the standardized deviation between adolescent rainfall and the long-run mean. Note that since all women of cohort c in village j have the same rainfall history, this is equivalent to a within-cluster cohort-weighted average of shocks, where the cluster-specific cohort shares ω_{cj} are the weights. Therefore, s_j represents the rainfall deviation experienced by the average female cohort in cluster j in their adolescent period. I average over all members of the women's DHS sample aged 15-49 in each village, though I test for differential cohort effects in Appendix F.

Standardizing the individual-specific shock s_{ij} accounts for both the mean level of rainfall in a cluster as well as its variability. This is germane given the literature on the effects of climate and geography on development (Dell et al., 2014), suggesting that that both the long-run level and variability of rainfall are correlated with present-day income. Since the shock is defined using rainfall data that often occurs decades (depending on the age-structure of the village) before the outbreak of Boko Haram violence, there is little danger of endogenous measurement error at weather stations as a direct result of conflict, as highlighted by Schultz and Mankin (2019).

5.3 Balance tests

I consider a first-pass test of the exogeneity of the instrument by running the primary reduced-form regression, including terms for current rainfall and its interaction with polygamy, substituting various confounders for the outcome variable. Under exogeneity of the excluded IV, this "balance test" should yield small and insignificant coefficients on the interaction term $p_{jr}s_{jr}$.

The top panel of Table B1 tests whether the interaction term is correlated with measures of cluster-level mean and dispersion of the wealth index. The mean wealth index score is uncorrelated with the differential adolescent rainfall shock. However, within-cluster wealth inequality, as measured by the Gini of wealth index scores, is significantly correlated with the instrument. If the rainfall shock affects inequality, then this may be an important channel through which shocks affect Boko Haram activity, violating the exclusion restriction. Therefore, I include wealth inequality among the controls in the estimation. In the second panel I consider the average age of men and women, the share of Muslims, the share of Hausa and Kanuri ethnicities, and the log of population density in 2005, as possible demographic confounds. For fixed characteristics, in panel 3, I consider the long-run (1901-2014) mean and SD of rainfall, the mean monthly temperature in the year of the survey, distance to water, distance to borders, altitude, slope, and latitude. Only average female age is significant at conventional levels.

In the final panel of Table B1, I consider two categories of non-Boko Haram conflict – riots and communal militia events – which are less likely to be driven by marriage, in contrast to Boko Haram, which explicitly appeals to the marriage market. These other forms of conflict are uncorrelated with the adolescent rainfall-polygamy interaction, suggesting that the IV does not affect all conflict, but only marriage-specific forms. It is reasonable to ask why these other forms of conflict do not respond, since the same mechanism might well be at play. Unlike with Boko Haram, there is no qualitative evidence that communal ethnic militias use marriage to recruit. Such militias are generally less organized and permanent, and do not exercise control over substantial territory. This suggests that they do not have access to the same recruitment technology, which requires substantial funding and organization to conduct raids, as well as control over territory to allow members to live freely with kidnapped brides.

6 Main results

The main results of the paper are presented in Table 2 and Table 3. In Table 2, I present results from the baseline specifications of the OLS, reduced-form, first-stage, and 2SLS regressions, which include the full set of controls as well as DHS round fixed effects, and are estimated for 3 outcome variables: Boko Haram incidence (1), number of events (2), and average annual fatalities.[20] In Table 3, I probe robustness of these baseline specifications to various combinations of controls and location fixed effects for Nigerian states and ethnic homelands.

6.1 OLS results

Consistent with Figure 2, the OLS results in Table 2, Panel A demonstrate a positive association between extremist violence and marriage inequality. This relationships holds across all three conflict outcomes: a one-unit increase in the marriage Gini is associated with a 0.3 percentage point increase in Boko Haram incidence, 0.042 additional events, and 0.48 additional deaths, with only the first outcome significant at conventional levels. Of course, marriage market concentration is likely to be correlated with wealth inequality, cultural factors, and a variety of omitted variables affecting conflict, making interpretations of the OLS results difficult. To make progress, we turn to the reduced form and instrumental variables results.

Table 2

6.2 Reduced form results

In the reduced form, we should see that Boko Haram conflict responds to adolescent rainfall shocks in polygamous places and is relatively unaffected in monogamous ones. To illustrate the identifying variation, I present results from the reduced form regression as a binned scatterplot in Figure 3, partialing out the full set of controls in the baseline specification. In Panel A, I run this regression only for clusters in which $p_{jr} = 1$; here we observe a strong correlation between the mean annual Boko Haram deaths and the magnitude of the adolescent rainfall shock. In Panel B, I estimate

[20] Incidence is a dummy variable indicating any Boko Haram activity in cluster, while number of events is the average annual number of Boko Haram-related events recorded in ACLED.

the same regression in the monogamous sample, and the positive correlation disappears.

Figure 3

To understand magnitudes and probe the robustness of these results, I estimate reduced form regressions in Panel B of Tables 2 and 3. In the baseline specification of Table 2, I estimate that a one standard deviation increase in rainfall above its mean during the average woman's adolescent period corresponds to a differential increase in polygamous villages of 27.5 percentage points in incidence,1.8 additional events, and 20.1 additional annual deaths, all significant at the 5% level. The estimates are very large, the latter being just less than 3 times the mean annual Boko Haram activity in the data.[21]

In Table 3, I test robustness to various combinations of controls and fixed effects. The reduced form estimate ranges from 8.1 to 20.1 deaths, depending on covariates. It is largest under ethnic-group fixed effects, both with and without controls, in columns (2) and (4). The inclusion of state fixed effects in columns (3) and (5) substantially reduces the effect sizes to 9.7 and 8.1, respectively. While this may be taken as a sign that the instrument is correlated with unobserved heterogeneity at the state-level[22], it is also possible that the fixed effects exacerbate measurement error bias in a small sample, which may be a particularly acute problem here given that the assignment of clusters to relatively coarse rainfall grids only imperfectly captures actual rainfall in a given cluster.

6.3 First stage results

The first-stage relationship estimates the regression of marriage inequality on $s_{jr} \times p_{jr}$, controlling for s_{jr}, p_{jr}, and a host of covariates. These models are of interest in their own right, as they provide a direct test of the theoretical mechanism – that income realizations in the pre-marital adolescent period should affect marriage decisions by shifting the reservation price, primarily in polygamous regions. This market-level first stage equation forms the basis of the IV strategy.[23]

To begin, I replicate the graphical analysis of the reduced form, plotting the (binned) relationship between the Gini coefficient of brides against the mean adolescent rainfall deviation in Figure 4. All variables are residualized for the full set of baseline controls and DHS round dummies. As

[21] However, it is important to keep in mind that since the deviations are averaged within clusters, a one-unit increase is equivalent to moving from roughly the 1st percentile of shocks to the 90th, as seen in Figure A3.

[22] This could be capturing differences in state-level policies, politics, geographic, or cultural features

[23] In Section 7, I estimate a set of individual-level regressions using age of marriage, child marriage, and the marriage hazard as dependent variables, which we can consider the "zero'th" stage, to demonstrate that this aggregate effect is driven by a micro-level marriage response.

before, Panel A shows the relationship in polygamous areas, which is positive and significant; in monogamous areas (Panel B), the slope is negative.

Figure 4

By raising reservation prices, supply-side income shocks should reduce marriage rates among poorer men and, given the correlation between income and brides, increase marriage inequality. As such, either unmarriage or inequality can be used as an endogenous variable in IV estimation, and they are highly correlated ($R^2 = 0.87$). I present first-stage estimates for both variables: in the main results of Table 2 and 3 Panel D, I use the Gini coefficient of brides. In the appendix Table B2, the bottom panel uses the unmarried rate among men, finding similar results.

In the baseline specification of Table 2 Panel D, a 1-unit differential increase in mean adolescent rainfall in polygamous areas is associated with a 12.8-point increase in the marriage Gini, nearly a 25% increase on the mean.[24]. Additional specifications in Table 3 do not materially affect the results. The estimate falls slightly with the addition of ethnic fixed effects in column (3), state fixed effects in column (4), or both in column (5). In the unconditional specification of column (1), the coefficient is largest, at 18.9. The pattern of results is nearly identical for the unmarried rate among men in Table B2. The results are not materially affected by the inclusion of controls, and attenuate slightly (by 4-5 percentage points) with the inclusion of state fixed effects. Across the preferred specifications in columns (2)-(5), a 1 SD increase in mean adolescent rainfall results in a differential (by polygamy) increase of 9.4-14.2 percentage points increase in the share of men that are unmarried, or roughly a 20.2-30.8% increase relative to the mean.

6.4 IV results

In this section, I estimate the effect of marriage market concentration on extremist violence via two-stage least squares, instrumenting the bride Gini with $s_{jr} \times p_{jr}$, conditioning on both s_{jr} and p_{jr}, current and lagged rainfall, as well as a bevy of controls and fixed effects. Under the assumption that $s_{jr} \times p_{jr}$ is conditionally mean independent of the error v_{jr}, 2SLS provides causal estimates of the parameter φ. This parameter measures the response of average annual fatalities by Boko Haram to a one-point (0.01-unit) increase in the Gini coefficient of brides in a village.

I present the baseline IV results in Table 2, Panel C, alongside the reduced-form, first-stage, and

[24] While this effect size appears large, note the distribution in Figure A3 Panel B shows that a one-unit increase covers most of the range. In Appendix D.6, I show similar effect sizes across a variety of inequality measures.

OLS results. Additional specifications are presented in Table 3. As before, the baseline model includes DHS round dummies and a full set of controls. For all IV results, I report the Kleibergen and Paap (2006) F-statistic of the first-stage regression, which accounts for clustering of error terms within rainfall-grid units. Column (1)-(3) indicate that a one-point increase in the bride Gini results in a 2 percentage point increase in conflict incidence (10% of the sample mean), an additional .14 events (19.8% of the sample mean), and 1.56 additional fatalities (20.5% of the sample mean). All estimates are significant at 5%. In the baseline specification of Table 2, the first stage F is 22.78, well above the Stock-Yogo threshold of 16.38 for a 10% maximal 2SLS size, obviating concerns about weak instruments. The IV estimate for total deaths (column (3)) is roughly 3 times larger than the OLS, the latter of which is significant only at the 10% level.

Table 3

The inclusion of fixed effects in Table 3 changes the result only slightly: 2SLS estimates range from 0.674 to 1.62 additional deaths per one-unit increase in the bride Gini. The result is smallest in the simplest specification (column 1), and largest when including both controls and ethnicity fixed effect. The inclusion of controls and fixed effects increases standard errors across the board, and also weakens the first stage. In the unconditional specification, the F-statistic is 36.18, but it falls to as low as 10.28 in the specification with controls and state fixed effects (column 5). In the full fixed effects specification, the IV estimate is 0.81 and the F is below the critical value for 10% maximal size, suggesting that limiting variation in shocks to within-state and within-tribe may introduce a weak instruments problem. However, to the extent that weak instruments bias the IV toward the smaller OLS (Angrist and Pischke, 2009) estimate, this makes finding an effect less likely.

Two key findings emerge from the analysis in Table 3. First, the IV effect is robustly large and significant: it ranges from an 8.8-21.3% increase in Boko Haram-related deaths relative to the mean annual number of deaths in a village (7.6), a large effect, but not implausibly so. Secondly, the IV results are substantially larger than the OLS results – between 2-3 times larger across all specifications. This discrepancy is likely to be driven by omitted regional characteristics leading to downward bias in the OLS estimates. A glance at the geographical distribution of polygamy in Figure A1 reveals that the polygamy rate is actually relatively low in Nigeria's northeast and high in other parts of the country. As the northeast is both poorer and more muslim – in addition to many other differences – we might see a downward bias in the cross-sectional correlation. This would also explain why the OLS estimate increases when we include controls and fixed effects in columns (2)-(5), relative to (1). It is also consistent with the fact that the OLS estimate falls to zero when we restrict to southern states, and rises to 0.64 when we restrict to northern states.[25]

The difference could also be driven by misreporting in self-reported marriage data or differences between LATE and ATE. If misreporting is noise and we assume the IV to be the true parameter, then a random measurement error story requires that the signal-to-noise ratio be very low to explain a 2-3-times magnitude difference. While there is evidence of misreporting of age (Lyons-Amos and Stones, 2017) and marital and reproductive activity (Neal and Hosegood, 2015) in DHS data, the magnitude seems implausible.[26] Still, OLS and IV estimates may still diverge if there are heterogeneous treatment effects for the sub-population affected by the instrument (Angrist and Imbens 1994, Angrist and Pischke 2009).[27] In our context, markets that have always been highly concentrated (as in the high-polygamy areas of Nigeria's south) may be less likely to fall into conflict than those in which high concentration is a state induced by transitory shocks. The sub-sample tests of Appendix Table D1 are suggestive of substantial heterogeneity: northern, rural, poor villages account for the average effects observed in Table 2.

In Appendix D.8, I formalize this discussion using a Bayesian estimation procedure developed by DiTraglia and Garcia-Jimeno (2018). In short, this procedure allows the researcher to estimate whether pre-specified beliefs on the sign of endogeneity and the extent of measurement error are consistent with a valid instrument and a causal effect of a given size. I consider whether the beliefs expressed informally above – negative selection and modest measurement error in marriage inequality – are consistent with the effect size estimated in Table 2 and an instrument uncorrelated with the error term u_{jr}. The results – presented in Table D8 and Figure D6 – generally support instrument validity and suggest that these beliefs are mutually compatible. For more detail, see Appendix D.8.

6.5 Age effects and marriage inequality

The reduced form and first stage results in Figures 3 and 4 both reveal the somewhat unexpected result that in monogamous areas, marriage inequality and less so Boko Haram conflict actually fall in response to the adolescent rainfall shock. Empirically, this is driven by a reallocation of brides across men of different age groups. Figure A5 plots the market share of men above and below 35 as a

[25] See Table D1.

[26] If survey respondents intuit that polygamy is perceived as socially undesirable by surveyors, they may under-report second or third wives, leading to an underestimate of marriage inequality in areas where it should be high, reducing the slope of the OLS relationship. Evidence suggests that populations may respond to surveyor identity in ways suggesting "demand effects" (Cilliers et al., 2015). Table B2 provides a test of this hypothesis by using the male unmarried rate – which will suffer less from such bias – as the endogenous variable. Here, the relative magnitudes of the OLS and IV estimate are similar to the main estimates.

[27] This is a common issue in the literature on returns to schooling (Oreopoulos, 2006).

function of the mean adolescent rainfall shock. When monogamous areas (Panels A and B) experience shocks, younger men gain market share, while older men lose; the opposite is true of polygamous areas (Panels C and D). Greater market share going to older men drives marriage inequality both because older men are more likely to have previous wives, and because they are richer, implying a higher bride price that squeezes young men out of the market. The result is greater Boko Haram activity, as shown in Figure A7.

Reallocation in marriage market share should be accompanied by changes in marital age gaps. This is confirmed by Figure A6, which shows spousal age gaps as a function of mean adolescent rainfall, by polygamy status. The cluster-level average spousal age gap falls dramatically in monogamous areas (Panel B), but remains flat in polygamous ones (Panel A). Note that it may not necessarily rise in polygamous areas largely because of the offsetting increase in age of marriage by women, as shown in Section 7.1. Table 4 tests the statistical significance of the spousal age gap effects under multiple specifications. Columns (1)-(3) demonstrate a null effect in the polygamous subsample, while columns (4)-(6) show a reduction in marital age difference ranging from 0.86 to 2.72 years, significant at 5% in 2 out of 3 specifications.

Table 4

Why do we observe reallocations in opposite directions along the age distribution? The theory presented in Section 3 assumes that preferences are defined only over income, so women can only react to income fluctuations by via the reservation price. However, if parents have preferences both over the age and bride price of their daughter's spouse, they may adjust to shocks along both margins.[28] A better bargaining position allows parents to pick spouses who are either richer (older) or more attractive (younger). In polygamous areas, the former effect dominates, since older men reenter the marriage market with high bride price offers. In monogamous areas, the latter effect may dominate, as the option to marry richer old men is diminished.

To illustrate this mechanism formally, consider the following extension of the original model. Now parents have preferences $u(a, y) = \alpha_1 \log(y) + \alpha_2 \log(1 - a)$, and men of type j are defined by the pair (a^j, w^j). We can think of the ratio $\frac{\alpha_1}{\alpha_2}$ as capturing the altruism of the parents toward the desire of their daughter to wed a younger man. For any $\alpha_2 > 0$, the qualitative predictions below will hold.

[28] Preference for youth may be driven by attractiveness, expected number of children, and/or expected investments in children. That is, older men will have fewer remaining years to have children, resulting in fewer children, as well as more extant children, resulting in less investment per child.

The parents' indifference condition for a man who offers (w, a) will be

$$\alpha_1 \log(w) + \alpha_2 \log(1-a) = \log(y^f) + \beta V_r^{wait} = \bar{u}_r$$

For simplicity, assume further that instead of discrete types, we have a unit measure of men of age a. Under polygamy, there is a positive correlation between wealth and age, since older, richer, married men are in the market, yielding an offer curve $w = \gamma a$. However, this correlation is absent under monogamy, as rich older men select out of the marriage market. Those who remain in the marriage market are either young men of any income type, or older men too poor to afford a wife. This assumption is well-founded in the data, as demonstrated in Appendix C.3. The indifference condition yields the relationship

$$w = y^f \exp(\beta V_r^{wait})(1-a)^{\frac{-\alpha_2}{\alpha_1}}$$

To analyze the decision rule of the family, Figure 5 plots the age-income profiles and the indifference curves in w, a space. All offers to the upper-left of reservation utility are accepted in equilibrium. The intersection between indifference curves and age-income profiles represent reservation points under different regimes and utility levels. The plot demonstrates that $i)$ inequality is still higher in polygamous areas, and $ii)$ positive income shocks raise the minimum age and bride price under polygamy, but reduce the maximum age and leave the bride price unchanged under monogamy.

First, we see that initial reservation utility $\bar{u}_p > \bar{u}_m$ because of differential continuation values of polygamy and monogamy. However, unlike in the initial analysis, this does not necessarily translate into higher bride prices. This is because offered bride prices vary by age under polygamy, so the average price must depend on this distribution. Still, under polygamy we are likely to see greater inequality, since it entails a minimum age cutoff, while there is no such cutoff under monogamy, where there is no compensating benefit to marrying older men.

Figure 5

A positive income shock increases reservation utility, shifting the indifference curve inward from $\bar{u}_m$ to $\bar{u}'_m$. Along the $w_m = c$ line, this reduces the maximum age a girl will accept from a_{max} to a'_{max}. This implies a falling market share for older men, reducing marriage inequality. In the polygamous market, a positive income shock also shifts the reservation utility from $\bar{u}_p$ to $\bar{u}'_p$. However, given the positive relationship between age and wives, this results in an increase in the minimum acceptable age from a_{min} to a'_{min}. Young men between these two points no longer find wives, since rising reservation utility requires that women demand higher bride prices. Given the convexity of preferences in this example, it is also possible that some older men are also shut out of the market. Still, the market share of younger men is likely to fall, increasing marriage inequality.

6.6 Threats to identification

The exclusion restriction identifying the model is that the interaction between the adolescent rainfall shock and the polygamy rate in cluster j is uncorrelated with u_{jr}, conditional on s_{jr} and p_{jr}, or that the differential (by polygamy rate) shock is not correlated with the unobserved determinants of current Boko Haram activity. In terms of the model, it must only affect the woman's decision by changing her income, leaving preferences and other marriage market features unaffected. While the identification strategy does not require polygamy itself to be exogenous, it does require that polygamous areas are not differentially affected by outside shocks that are correlated with adolescent rainfall.

In this section, I address two important plausible violations. Firstly, if weather shocks are serially correlated, the adolescent rainfall shock might be correlated with current rainfall. This is important given the well-established relationship between current weather shocks and political conflict (Harari and Ferrara 2018, von Uexkull et al. 2016, Burke et al. 2015). Even in the absence of serial correlation, if shocks have persistent effects (as in Maccini and Yang 2009), then the adolescent rainfall shocks could affect current outcomes through non-marriage channels. Secondly, polygamy may be proxying for other cultural norms or ethnoreligious differences, and these may also interact with adolescent rainfall shocks. An obvious example is Islam, which exhibits a strong unconditional correlation with polygamy ($R^2 = 0.23$) at the cluster level. Islam, in turn, is also correlated with ethnicity, as substantial differences in religious practice exist across ethnic groups in Nigeria. If any of the ethnicity or religion-specific characteristics correlated with polygamy also mediate the differential response to adolescent rainfall shocks, the interaction term IV may be picking up these effects.

6.6.1 Autocorrelation in weather

The exclusion restriction is violated if there is autocorrelation in weather shocks, or if income effects of past shocks materialize with a lag, so that shocks during the adolescent period of the cluster's average woman directly impact either economic conditions or conflict today. There are three reasons why this is unlikely: $i)$ this effect must be differential with respect to polygamy, as the identification strategy differences out the level effects, $ii)$ direct income effects would have to operate with a substantial lag or a high degree of autocorrelation given average age of women in the sample (28.7), and *iii*), the most comprehensive evidence suggests that good rainfall should reduce, rather than increase, the incidence of conflict (e.g. Harari and Ferrara 2018); however, we observe the opposite.

To address autocorrelation, I control directly for current and one-period lagged rainfall and their interaction with polygamy; if autocorrelation is not a concern, this should not materially affect the estimates of π_1 or φ. Table B3 shows the results for the first stage (columns (1)-(4)) and reduced form (columns (5)-(8)) regressions. In columns (1) and (5), I reprint the baseline results without controlling for the rainfall interactions. Including current and lagged rainfall deviation and their interaction with the polygamy dummy does not materially change the first stage coefficient, and produces similar reduced form estimates to those in Table 3. Polygamous areas do not respond significantly differentially to current shocks along either outcome, so that the assumption of no differential non-marriage effects of rainfall is plausible. Table B4 re-estimates the 2SLS models. By increasing noise, the current-rainfall interaction weakens the first stage enough to fail the Stock-Yogo 10% size test in some specifications. Therefore, for all significance tests I report weak-IV robust Anderson-Rubin p-values. The IV results are not much changed from those in Table 3; though the AR p-values indicate that only 3 of the 8 robustness tests remain significant at 5%, the estimated φ do not change.

To rule out persistent non-marriage income effects of past shocks, I consider the following placebo test. For each woman I construct rainfall shocks that cover the period beginning 8 years before her birth year until 17 years after birth. Shocks are constructed identically to s_{ij} except that they use different four-year windows for rainfall deviation. For each shock s_{ij}^{τ} corresponding to an initial year of the four-year window relative to birth, $\tau \in [-8, 17]$, I estimate the first stage, and the reduced form. I then plot a histogram of estimated coefficients with a vertical line indicating the value corresponding to $\tau = 12$ (Panel A), and well as a plot of their time path relative to birth with 95% confidence intervals (Panel B). The coefficient of the true adolescent rainfall shock should be in the right-tail of the distribution of estimates obtained from these placebo tests.[29] We also expect the time-path of coefficients to show a monotonic increase over time until the adolescent period, peaking near (but not necessarily at) $\tau = 12$. Qualitatively, these patterns emerge in the analysis of equilibrium cluster-level marriage outcomes in Figure A8 and in the reduced form conflict equation in Figure A9 as well.

6.6.2 Ethnicity and religion

Polygamy is correlated with specific cultural features. It is a more important institution in areas that have larger muslim population shares, as well as among the dominant ethnic group in Northern

[29] Note that it may not necessarily be the largest; similarly-sized effects may be obtained by shocks around adolescence but in different windows (i.e. ages 13-17) given that 12-16 was an arbitrary choice.

Nigeria, the Hausa.[30] This is a problem if, for example, good adolescent rainfall may relieve credit constraints and lead to increased investment in girls' education, an effect which could depend on religion or ethnicity. This could have separate knock-on effects on both marriage markets and conflict in ways that are unrelated to the explanations proposed here. Throughout the paper, I include the levels of these demographic variables in the main set of controls X_{jr}. As an additional test, I re-estimate the reduced-form and first-stage equations with the interaction term $s_{jr} \times d_{jr}$, where d_{jr} is a given demographic population share in village j at DHS round r.

The results are presented in Table B5; the top panel re-estimates the reduced-form, while the bottom re-estimates the first stage. Columns (1)-(4) are the unconditional specification, while columns (5)-(8) include ethnic homeland fixed effects as an additional robustness test. Columns (1) and (5) reprint the main results under either specification for reference. In columns (2) and (6), including the interaction with Kanuri share, has no effect on the first stage. Columns (3) and (7) add the interaction with Muslim share, which again has no effect. However, including the Hausa share attenuates the first stage to from 12.9 to 11.2 without fixed effects and from 11.8 to 9.2 with them. When we consider the reduced form (top panel), the inclusion of the Kanuri share still has no effect on the estimated coefficient. However, the inclusion of interactions with the Muslim share reduces the estimate from 20.6 to 10.9 under no fixed effects (column (3)) and reduces the estimate from 18.4 to 11.8 (column (7)) under ethnic homeland fixed effects. They are reduced even further, to 8.6 and 7.7, respectively when we consider d_{jr} to be the Hausa population share. However, the reduced form remains significant in all specifications except column (4).

6.7 Additional robustness tests

In Appendix D, I explore the robustness of results to sample restrictions (D.1), variable definitions (D.2, D.6), data sources (D.3), IV diagnostics (D.8), migration (D.4, D.5), demand-side effects (D.7), age effects (D.11), marriage market size (D.9), types of violence (D.10), and IV inference (D.12). Lastly, in Appendix F, I conduct an additional falsification exercise which demonstrates that the pattern of effects across cohorts follows what would be expected if a positive income shock permanently reduced the per-period marriage hazard over the lifetime. A few tests are worth highlighting.

Demand side effects: The exclusion restriction is violated if the female adolescent rainfall shocks s_{jr} also capture positive income effects for young men that have direct effects on the marriage market or subsequent conflict. The identification assumption benefits from the fact that while there are large

[30] This is true in both in bivariate and conditional correlations. Results available upon request.

gender gaps in age of marriage, the age structure in symmetric across genders. Female adolescent rainfall shocks may well be highly correlated with the equivalent male adolescent rainfall shocks.[31] However, male shocks should not affect men's marriage market prospects, since 12-16 are not the pre-marriage market years for men, who marry on average a decade later at age 25. To test this, in Appendix D.7 I include male-side shocks for various age ranges interacted with polygamy in the reduced form, first stage, and 2SLS regressions. The magnitude of the results are not much affected, though the correlation between the male and female shocks substantially increases standard errors.

Recruitment assumption: The model makes predictions on how the joining rate ρ will respond to exogenous shocks and marriage market institutions. However, only data on attacks and fatalities are available. In Appendix D.10, I test indirectly the validity of the assumption that violence is a good proxy for recruitment, exploiting information contained in the spatial variation of different types of attacks. Using previously conducted small-scale surveys of ex-combatants, I identify the states that account for the majority of Boko Haram recruitment. I then use event-level data to identify types of violence – suicide attacks, bombings, and remote violence – that are more often committed in areas distant from Boko Haram's core recruitment areas. I then re-estimate the main models dropping these high mobility attacks which are likely to be poor proxies for recruitment. The results hold.

Marriage market size: The empirical analysis also assumes that a village is an appropriate marriage market unit. Of course, there may be migration for marriage, and markets may be spatially integrated. To the extent that young men can search outside their home villages, shocks in a given village will propagate to its neighbors, as men facing higher bride prices exit the village. These spillovers would moderate the impact of supply shocks on bride prices, likely leading me to underestimate the true reduced form and first stage effects. In Appendix D.9, I test robustness to this source of bias by aggregating to larger marriage market clusters using a k-means distance-based clustering algorithm. The results do not change, though the significance declines because of the smaller sample.

6.8 Falsification test: Non-Boko Haram violence

The patterns of non-islamist violence in Nigeria provides an additional falsification test to support the proposed causal mechanism. Most non-islamist violence in Nigeria is political or ethnic in nature, but does not take the form of organized insurgency. Community-based ethnic militias, politically-affiliated militia groups, and criminal gangs – responsible for 52% of Nigeria's non-Boko Haram ACLED events – are not able to exploit marriage market incentives as effectively as Boko Haram.

[31] Indeed this is the case, with an $R^2 = 0.64$.

This is primarily because they lack the organization, scale, and control of territory required to do so. Organization and scale are necessary to carry out the strategic raids that are required to satisfy the demand for captive brides among prospective recruits. Administrative control of territory is necessary to be able to facilitate mass weddings and formalize marriages so that fighters may live unmolested with captive brides in areas governed by the rebel group. Thus, to exploit marriage market imbalances, a group must be able to carry out large-scale raids and control territory. During the period of study, no armed group in Nigeria other than Boko Haram satisfies these criteria.

We should therefore observe that non-Boko Haram violence does not respond to the changing marriage market induced by adolescent rainfall shocks and marriage market institutions. I show in Table B6 that this is indeed the case. Using either non-Boko Haram attacks in columns (1)-(3) or non-Boko Haram fatalities in columns (4)-(6), there is no significant effect of the interaction between adolescent rainfall and polygamy. Even the OLS correlations are very close to zero. These null results are not sensitive to specification, and are precisely estimated relative to the main results. It is clear that only Boko Haram are capable of exploiting marriage market incentives.

7 Mechanisms

In this section, I test several hypotheses implied by the theoretical model to support the underlying mechanism driving the results presented in Section 6. Firstly, the aggregate, market-level first-stage relationship should be driven by changes in individual marriage choices: in response to positive adolescent shocks, women should reduce their marriage hazard, marry later on average, and marry the richer and more polygamous men who can afford higher bride prices. This selection effect should underly the increase in inequality observed at the market level. Consequently, there should also be an increase in bride prices corresponding to the higher willingness to pay of selected men. Finally, the number of abductions should increase, since more men are willing to join Boko Haram so the supply of captured brides must expand.

7.1 Individual-level marriage results

Do aggregate results reflect an underlying supply response at the individual level? I test for this response using the individual sample of women in the DHS data. To begin, I regress the reported age of first marriage on the average annual (standardized) deviation of rainfall from its long-run mean in the years when a woman is aged 12-16. I also consider child marriage as an outcome by regressing an indicator for marriage on the adolescent rainfall deviation in the sample of women aged 18 or less.

The estimating equation for woman i in grid-cell j in cohort c is

$$y_{i,j,c} = \alpha + \sigma s_{i,j,c} + \xi_j + \psi_c + \delta_r + X'_{i,j,c}\beta + \eta_{i,j,c}$$

Where $y_{i,j,c}$ is the marriage outcome of interest and the controls $X_{i,j,c}$ include a quadratic polynomial in age, dummies for Muslim and Hausa, and current rainfall deviation. For the age of marriage analysis, I estimate this model using both OLS as well as a two-step Heckman selection procedure to account for selection into the ever-married sample. As selection variables I use age, age-squared, years of education, year of birth, and indicators for rural, muslim, Hausa, and Kanuri. Table 5 contains estimates for these individual-level specifications. Columns (1)-(4) estimate the model with age of marriage as y, while columns (5)-(8) use child marriage. Panel A presents results for OLS, while Panel B presents selection models. In the individual-level data, the effect of interest σ is identified by rainfall variation across locations for a given cohort, and rainfall variation across cohorts for a given location, allowing fixed effects for both rainfall grid-cell ξ_j and cohort ψ_c. These capture unobserved, cohort and location specific heterogeneity that might correlate with adolescent rainfall shocks.

Table 5

The results of Table 5 show that the supply response to economic shocks operates at the individual level. In the preferred specification of column (4), one SD increase in average annual rainfall over the adolescent period above the long-run mean is associated with 0.542 years delayed marriage on average. This estimate is stable once we account for local heterogeneity, with state (3) and grid-cell (4) fixed effects producing similar estimates. The estimate is nearly doubled in size when we omit geographic fixed effects entirely in (1) and (2), suggesting that geographic heterogeneity biases the coefficient upwards. This seems reasonable if locations more likely to experience positive rainfall shocks are different, but the timing of those rainfall shocks across cohorts within a location is random. A similar set of results obtains for child marriage. Accounting for selection in Panel B does not materially change the results in columns (1)-(2), but actually leads to substantially larger estimates in columns (3) and (4), implying a downward bias induced by sample selection. The coefficient on λ, the inverse-Mills ratio, is positive and significant, rejecting the null of no selection. In columns (5)-(8), a one SD increase in adolescent rainfall reduces the likelihood of a marriage among girls below the age of 18 by 3.9-7.1 percentage points. The results are robust to the inclusion of the controls and fixed effects, although they lose statistical significance in the most rigorous specification (column (8)).

In Figure 6, I plot estimates of σ using pre-and-post adolescent shocks s^{τ}_{ijc}. In Panel B, the time path of coefficients is inconsistent with persistent spurious income effects. Rainfall shocks prior to birth and in the very early years of childhood have essentially no effect on the marriage decision. The estimated coefficients for $\tau < 4$ are all near negative or zero, and mostly insignificant. After age four, the importance of shocks increases dramatically as parents likely begin to think about marriage for their daughters. By $\tau = 7$, a 1 SD shock increases age of marriage by about 0.2, doubling by $\tau = 10$. The effect size appears to peak slightly after our chosen period, at $\tau = 14$.[32]

Figure 6

In a final set of supply-side tests, I estimate the hazard into marriage as a function of adolescent rainfall, as in Corno et al. (2017). Using women's age of marriage, I form an individual-level panel which contains for a woman i in calendar year t the annual rainfall in grid j, her age a, cohort c, and marital status dummy m. Each woman is retained in the data from age ten until she exits the sample via marriage or until age 45. To test the effect of above-average rainfall in the adolescent period on the hazard rate into marriage over the lifecycle, I estimate the following regression

$$m_{i,j,a,c,t} = \alpha + \eta_0 s_{j,c} + \sum_{k=11}^{45} s_{j,c}\eta_k 1(a_{i,t} = k) + \zeta_a + \gamma_j + \delta_c + \mu_{i,j,a,c,t}$$

The specification allows the effect of the adolescent rainfall shock $s_{i,j,c}$ to vary with age. The regression also contains fixed effects for location, cohort, and age. I estimate the model separately by polygamy status to test the contrasting prediction of the theory–that the hazard rate should fall as income rises, and that these effects should be more pronounced in polygamous areas.[33]

Figure 7

I plot the η_k coefficients in Figure 7 for polygamous (Panel A) and monogamous (Panel B) areas separately. All estimates are the effect adolescent rainfall on marriage hazard at age a relative to

[32] This suggests I could improve the power of the first stage by shifting the adolescent rainfall shock period by two years. However, I use a pre-specified, arbitrary threshold taken from prior literature to allay concerns about p-hacking.

[33] Note that I use the interaction of adolescent rainfall and age to generate life-cycle variation of the adolescent rainfall shock within individuals. This is in contrast to a specification using simply the current-year rainfall deviation, which seems a more direct implementation of the theoretical mode, as reservation prices depend on current income. However, I use adolescent rainfall in the hazard model in order to use the same underlying variation as in the market-level equation. Theoretically, there are several justifications for this specification. If reservation wages are updated with a lag, if they respond asymmetrically to positive and negative shocks, or if timing matters, then rainfall upon market entry should more accurately capture the economic circumstances relevant to the formation of reservation prices.

the effect at age 10. The vertical line indicates age 16, the age of entry into the marriage market. The life-cycle pattern of responses to shocks confirms the predictions of the theory. In polygamous areas, good adolescent rainfall reduces hazard into marriage substantially; reassuringly this effect only emerges after age 16. This effect grows and peaks during the mid-20s, after which it attenuates to zero as around age 34. In monogamous areas, the hazard effect is initially negative and significant after age 16, though small in magnitude. However, this relationship exhibits no clear pattern in the years after age 25, and is not statistically different from zero.

7.2 Match quality

Increasing concentration implies that in times of plenty, women should match with higher quality (richer) men, who are also more likely to be polygamous.[34] This latter prediction is not entirely obvious; an alternative hypothesis holds that if women incur disutility from polygamy, a better bargaining position might reduce the selection into polygamous households, which might suggest less marriage market concentration.[35]

I use the individual-level cross-sectional specification, regressing husband type on adolescent rainfall deviation, individual-level controls, and cohort and state or grid-cell fixed effects. To measure how the expected husband type varies with adolescent rainfall, I use the spousal wealth index as the dependent variable.[36] In order to test whether women select into polygamous marriages as a result of the shock, I use an indicator for whether the husband was already married at the time of marriage as the dependent variable. For each outcome, I split the sample into the polygamous and monogamous clusters and estimate on these subsamples.

Table 6

The results of Table 6 confirm that women who experience good adolescent shocks marry richer polygamous men, driving the increase in market-level concentration. Columns (1)-(4) estimate expected type, while columns (5)-(8) estimate the likelihood of matching with a polygamous husband. In columns (1) and (2) I estimate that women who experienced 1 a SD positive rainfall deviation have a husband roughly 0.03-0.08 points richer on the 5 point wealth scale, significant at 5 or 1%. The effect vanishes in monogamous areas, consistent with the results throughout the paper that have

[34] This follows from the fact that richer men have a greater propensity toward polygamy, as seen in Table C2.

[35] This disutility is omitted from the model.

[36] Strictly speaking the variable available in the DHS data is post-marital household wealth. However, the results hold in households where women do not work, thus ruling out that effects are driven by the female contribution to household wealth. These are available upon request.

shown negligible effects of the shock on outcomes in these villages. Similar patterns are evident for the polygamous at marriage indicator. Women who experience a 1SD adolescent shock are 1.1-1.5 percentage points more likely to marry into a polygamous household in polygamous regions, but no more likely to in monogamous areas,[37] though the effect is only significant at 10%.

7.3 Marriage expenditures

Although the model predictions about bride price are ambiguous, evidence that prices are differentially affected by shocks in polygamous areas would support the proposed causal mechanism. In lieu of bride prices, I use data on marriage ceremony expenditures from Nigeria's General Household Survey (GHS), a nationally representative panel survey spanning 3 bi-annual waves from 2010-2015. This data has two important drawbacks: i) these are only observed for households that contributed to some marriage in the 12 months prior to the survey, and ii) these expenditures are likely to capture bride prices only imperfectly.[38] Unless the measurement error induced by over-estimating bride prices by including additional expenditures is systematically correlated with the variation of interest, it should not bias the results.

I test whether wedding payments increase in positive adolescent income shocks by regressing the log of the amount paid at the individual level on the cluster-level mean adolescent rainfall shock (constructed identically as in the DHS data), year dummies, state dummies, and a controls for rainfall, lagged rainfall, share muslim, distance to road, and average age of women in the village. I estimate this model on the polygamous and monogamous subsamples, using the same criteria to define the presence of a polygamy norm as throughout the paper. All specifications pool all available observations of marriage spending across all three panel waves, use combined sample weights, and cluster standard errors at the rainfall-grid level.

Table 7

The results are given in Table 7. The split-sample approach reveals substantial differences by polygamy status.[39] A 1 SD positive adolescent rainfall shock increases marriage expenditures by roughly 71-92% (columns (4)-(6)) under polygamy while there is no significant effect among monogamous villages (columns (7)-(9)) once we account for state-level heterogeneity. I also replicate the

[37] This latter result is essentially by construction.

[38] However, in the Nigerian context, marriage costs as reported in a survey are likely to be a reasonably good proxy for bride prices, since bride wealth is typically transferred by the groom around the time of marriage. When asked about wedding costs, men incorporate such transfers to the family of bride in these costs.

[39] In an interaction term specification, the heterogeneity is not significant since we are underpowered for differential effects.

placebo test to additionally support the identification assumption, using only the sample of polygamous villages, since this is where we observe the positive effects. To reduce noise in the estimation, I omit all controls except 2010 rainfall deviation and mean female age.[40] Figure 8 confirms that rainfall shocks experienced in the years between 8 years before birth and roughly 4-5 years after birth are uncorrelated with present marriage expenditures, while those experienced after this period have substantial and increasing positive effects.

Figure 8

7.4 Abductions and violence against women

When marriage market conditions worsen in polygamous regions as a result of positive premarital economic shocks for women, men join Boko Haram. To meet growing demand in this non-civilian marriage market, Boko Haram must procure more brides. In addition to attention-grabbing mass abductions, the group also conducts frequent abductions of both women and men on a smaller scale. Such abductions appear to be the primary way in which Boko Haram is able to obtain brides for its troops.[41] Empirically, we should observe an increase in abductions in times and places when the marriage market is exogenously tightened because of the interaction between economic shocks and the practice of polygamy. We should also observe an increase in violence against women, who are likely to be disproportionately targeted by rising violence. These effects should be over and above the response of overall conflict to these shocks.

In total, abductions occur in 3.27% of the DHS clusters, while violence against women occurs in 4.04% a relatively small proportion of the overall data.[42] There is a legitimate concern that, in some cases, these abductions do not refer to the abductions of women. If we restrict the sample further to consider only events that contain the kidnapping strings and a reference to women, such events occur in less than 1% of the clusters. This is far too small to be useful for analysis, so I consider separately the full sample of abductions and the full sample of attacks against women – defined as any violent Boko Haram event that includes reference to women or girls. There is also a concern that the increasing abductions might violate the exclusion restriction, since this directly affects the supply of brides in the civilian sector. I do not have the cluster-level data on the number of marriage-age

[40] Similar patterns are obtained when I include the full set of controls and state fixed effects, although the confidence bounds of the estimates are wider and the pattern is less monotonic.

[41] However, as Matfess (2017) makes clear, the abduction narrative is overly simplistic. Women also willingly join the group, though the frequency is unknown.

[42] For reference, any Boko Haram attack occurs in 22.13% of the DHS clusters. Given that the variable is zero in most clusters, floor effects imply that we should be unlikely to find an effect even if one exists.

women abducted required to rule this out. However, the extreme rarity of abduction events suggests that this effect would be small and can safely be ignored.

Table 8

Table 8 provides the results, which given the measurement caveat, should be cautiously interpreted. Columns (1)-(3) use an abductions dummy as the outcome variable, while (4)-(6) use a violence against women dummy. The results are consistent with the hypothesis that adolescent rainfall shocks increase (differentially by polygamy) these particular types of violence, even conditional on overall militant activity. The results on abductions are positive and significant in all specifications, while those on violence against women are also positive but not always significant.

8 Concluding remarks

It is well known that societies can be destabilized when a large mass of young men are excluded from the marriage market. This paper provides new evidence that traditional marriage practices can exacerbate marriage market tightness, highlighting a novel economic rationale underlying the extreme gender-based violence of insurgent groups like Boko Haram. When young women experience good income realizations in the years before they enter the marriage market, they raise their standards over the types of men they are willing to marry. This is manifested in lower annual marriage hazard rates, increased age of marriage, higher average bride prices, and selection into richer and more polygamous households. Critically, this effect is amplified by the existence of a polygamy norm. Where polygamy is widely practiced, women have the option to draw from a pool of richer, older, already-married men. The existence of this option increases the bride price women must be offered, and also their sensitivity to changes in their outside option.

These choices affect marriage markets, with severe social consequences. Young men in polygamous markets where women have good outside options face a choice: enter the civilian labor and marriage market with a low expectation of success, or turn to insurgents, who offer the promise of marriage and income. The more that women raise their standards because of polygamy and outside options, the more appealing the insurgent choice becomes. This leads to greater recruitment into insurgent ranks and, since insurgent strength depends on the size of their force, more overall violence. Consistent with this, polygamous villages where the average woman experienced good rainfall in her adolescence have lower marriage rates, higher marriage inequality, and more Boko Haram violence.

Insurgent groups also respond strategically; as unmarried men become more plentiful, abductions and violence against women are used more intensively to satisfy demand.

Standard theory and existing evidence suggests that improved economic conditions increase the opportunity costs of violence. These findings ignore the role of marriage market imbalances in promoting violence. Once this mechanism is accounted for, good economic conditions may actually exacerbate violence depending on the timing of shocks and prevailing marriage norms, though the underlying logic of opportunity costs remains intact. Nor is the Northern Nigerian case particularly exceptional. Many traditional, patriarchal societies are structurally similar, exhibiting polygamy, bride price, large age differentials at marriage, marriage inequality, and violence against women. Therefore, the results should be seen as plausibly generalizable for societies with similar cultural traits and marriage market conditions. Taken together, the results highlight the marriage market as a critical mechanism linking economic shocks to violence, and complicate the standard opportunity cost paradigm that has dominated the empirical study of conflict.

References

Anderson, Siwan, "The Economics of Dowry and Brideprice," *Journal of Economic Perspectives*, 2007, *21* (4), 151–174.

Angrist, Joshua D. and Adriana D. Kugler, "Rural Windfall or a New Resource Curse? Coca, Income, and Civil Conflict in Colombia," *The Review of Economics and Statistics*, 2008, *90* (2), 191–215.

_ **and Guido W. Imbens**, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, *62* (2), 467–475.

_ **and Jorn-Steffen Pischke**, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton: Princeton University Press, 2009.

Arthi, Vellore and James Fenske, "Polygamy and child mortality: Historical and modern evidence from Nigeria's Igbo," *Review of Economics of the Household*, 2018, *16*, 97–141.

Ashraf, Nava, Natalie Bau, Nathan Nunn, and Alessandra Voena, "Bride Price and Female Education," *NBER Working Paper No. 22417*, 2016.

Bazzi, Samuel and Christopher Blattman, "Economic Shocks and Conflict: Evidence from Commodity Prices," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 1–38.

Becker, Gary S., "A Theory of Marriage: Part II," *Journal of Political Economy*, 1974, *82* (2), S11–S26.

—, *A Treatise on the Family*, Cambridge: Harvard University Press, 1981.

Blattman, Christopher and Jeannie Annan, "Can Employment Reduce Lawlessness and Rebellion? A Field Experiment with High-Risk Men in a Fragile State," *Economics Letters*, 2016, *110* (1), 1–18.

Burke, Marshall, Solomon Hsiang, and Edward Miguel, "Climate and Conflict," *Annual Review of Economics*, 2015, *7*, 577–617.

Cameron, Lisa, Xin Meng, and Dandan Zhang, "China's Sex Ratio and Crime: Behavioural Change or Financial Necessity?," *The Economic Journal*, 2017, *Forthcoming* (1).

Cilliers, Jacobus, Oeindrila Dube, and Bilal Siddiqi, "The white-man effect: How foreigner presence affects behavior in experiments," *Journal of Economic Behavior and Organization*, 2015, *118*, 397–414.

Cohen, Ronald, "Marriage Instability among the Kanuri of Northern Nigeria," *American Anthropologist*, 1961, *63*, 1231–1249.

Cold-Ravnkilde, Signe and Sine Plambech, "Boko Haram: From local grievances to violent insurgency," Technical Report DIIS Report No. 2015:21, Danish Institute for International Studies 2015.

Corno, Lucia and Alessandra Voena, "Selling daughters: age of marriage, income shocks and the bride price tradition," *IFS Working Paper W16/08*, 2016.

—, Nicole Hildebrandt, and Alessandra Voena, "Age of Marriage, Weather Shocks, and the Direction of Marriage Payments," *NBER Working Paper No. 23604*, 2017.

Davidson, Russel and James G. MacKinnon, "Wild Bootstrap Tests for IV Regression," *Journal of Business and Economic Statistics*, 2010, *28* (1), 128–144.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, "What Do We Learn from the Weather? The New Climate Economy Literature," *Journal of Economic Literature*, 2014, *52* (3), 740–798.

DiTraglia, Francis and Camilo Garcia-Jimeno, "A Framework for Eliciting, Incorporating, and Disciplining Identification Beliefs in Linear Models," *Working Paper*, 2018.

Donnay, Karsten, Eric T. Dunford, Erin C. McGrath, David Backer, and David E. Cunningham, "Integrating Conflict Event Data," *Journal of Conflict Resolution*, 2018, pp. 1–28.

Dube, Oeindrila and Juan F. Vargas, "Commodity Price Shocks and Civil Conflict: Evidence from Colombia," *The Review of Economic Studies*, 2013, *80* (4), 1384–1421.

Eck, Kristine, "In data we trust? A comparison of UCDP GED and ACLED conflict events datasets," *Cooperation and Conflict*, 2012, *47* (1).

Edlund, Lena, Hongbin Li, Junjian Yi, and Junsen Zhang, "Sex Ratios and Crime: Evidence from China," *The Review of Economics and Statistics*, 2013, *95* (5), 1520–1534.

Erulkar, Annabel S. and Mairo Bello, "The Experience of Married Adolescent Girls in Northern Nigeria," Technical Report, The Population Council 2007.

Fenkse, James, "African polygamy: Past and present," *Working Paper*, 2013.

_ , "African polygamy: Past and present," *Journal of Development Economics*, 2015, *117* (November), 58–73.

Gehring, Kai, Sarah Langlotz, and Stefan Kienberger, "Stimulant or depressant? Resource-related income shocks and conflict," *University of Heidelberg Economics Discussion Paper Series No. 652*, 2018.

Grossbard, Amyra Shechtman, "The economics of polygamy," *Research in Population Economics*, 1980, *2*, 321–350.

Grossbard, Shosana, "An Economist's Perspective on Polygyny," in Janet Bennion and Lisa Fishbayn, eds., *The Polygamy Question*, 1 ed., University Press of Colorado, 2015, chapter 5, pp. 103–114.

Harari, Mariaflavia and Eliana La Ferrara, "Conflict, Climate and Cells: A Disaggregated Analysis," *Review of Economics and Statistics*, 2018, *Forthcoming.*

Hodler, Roland and Paul A. Raschky, "Economic shocks and civil conflict at the regional level," *Economics Letters*, 2014, *124* (3), 530–533.

Hudson, Valerie M. and Andrea M. den Boer, *Bare Branches: The Security Implications of Asia's Surplus Male Population*, Cambridge: MIT Press, 2004.

_ **and Hilary Matfess**, "In Plain Sight: The Neglected Linkage between Brideprice and Violent Conflict," *International Security*, 2017, *42*, 7–40.

Humphreys, Macartan and Jeremy Weinstein, "Who Fights? The Determinants of Participation in Civil War," *American Journal of Political Science*, 2008, *52* (2), 436–455.

Kanazawa, Satoshi, "Evolutionary Psychological Foundations of Civil Wars," *The Journal of Politics*, 2009, *71* (1), 25–34.

Kleibergen, Frank and Richard Paap, "Generalized reduced rank tests using the singular value decomposition," *Journal of Econometrics*, 2006, *133* (1), 97–126.

Lowes, Sara and Nathan Nunn, "Bride price and the wellbeing of women," *UNU-WIDER Working Paper 2017/131*, 2016.

Lyons-Amos, Mark and Tara Stones, "Trends in Demographic and Health Survey data quality: an analysis of age heaping over time in 34 countries in Sub Saharan Africa between 1987 and 2015," *BMC Research Notes*, 2017, *10*, 760.

Maccini, Sharon and Dean Yang, "Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall," *American Economic Review*, 2009, *99* (3), 1006–1026.

Matfess, Hilary, *Women and the War on Boko Haram: Wives, Weapons, Witnesses*, Chicago: University of Chicago Press, 2017.

Mberu, Blessing, "Who moves and who stays? Rural out-migration in Nigeria," *Journal of Population Research*, 2005, *22* (2), 141–161.

McGuirk, Eoin and Marshall Burke, "The Economic Origins of Conflict in Africa," *NBER Working Paper No. 23056*, 2017.

Miguel, Edward, Shanker Satyanath, and Ernest Sergenti, "Economic Shocks and Civil Conflict: An Instrumental Variables Approach," *Journal of Political Economy*, 2004, *112* (4), 725–753.

Moscona, Jacob, Nathan Nunn, and James A. Robinson, "Social Structure and Conflict: Evidence from Sub-Saharan Africa," *NBER Working Paper No. 24209*, 2018.

Murdock, George P., *Ethnographic Atlas*, Pittsburgh: University Of Pittsburgh Press, 1967.

Neal, Sarah E. and Vicky Hosegood, "How Reliable Are Reports of Early Adolescent Reproductive and Sexual Health Events in Demographic and Health Surveys?," *International Perspectives on Sexual and Reproductive Health*, 2015, *41* (4), 210–217.

Nunn, Nathan, "The Long-Term Effects of Africa's Slave Trades," *Quarterly Journal of Economics*, 2008, *123* (1), 139–176.

— **and Nancy Qian**, "US Food Aid and Civil Conflict," *American Economic Review*, 2014, *104* (6), 1630–1666.

Onuoha, Freedom C., "Why Do Youth Join Boko Haram?," Technical Report Special Report 348, United States Institute of Peace, 2301 Constitution Ave., NW, Washington, DC 2014.

Oreopoulos, Phillip, "Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter," *American Economic Review*, 2006, *96* (1), 152–175.

Raleigh, Clionadh, Andrew Linke, Havard Hegre, and Joakim Karlsen, "Introducing ACLED: An Armed Conflict Location and Event Dataset," *Journal of Peace Research*, 2010, *47* (5).

Sarsons, Heather, "Rainfall and Conflict: A Cautionary Tale," *Journal of Development Economics*, 2015, *115* (July), 62–72.

Schultz, Kenneth A. and Justin S. Mankin, "Is Temperature Exogenous? The Impact of Civil Conflict on the Instrumental Climate Record in Sub?Saharan Africa," *American Journal of Political Science*, 2019, *Forthcoming.*

UNDP, "Journey to Extremism in Africa," Technical Report, The United Nations Development Programme, 1 UN Plaza, New York, NY 10071, USA 2017.

von Uexkull, Nina, Mihai Croicu, Hanne Fjelde, and Halvard Buhaug, "Civil conflict sensitivity to growing-season drought," *PNAS*, 2016, *113* (44), 12391–12396.

Weeraratne, Suranjan, "Theorizing the Expansion of the Boko Haram Insurgency in Nigeria," *Terrorism and Political Violence*, 2017, *29* (4), 610–634.

Weidmann, Nils B., "On the Accuracy of Media-based Conflict Event Data," *Journal of Conflict Resolution*, 2015, *59* (6), 1129–1149.

Wu, Chien-Fu Jeff, "Jackknife, bootstrap and other resampling methods in regression analysis," *Annals of Statistics*, 1986, *14*, 1261–1350.

Young, Alwyn, "Consistency without Inference: Instrumental Variables in Practical Application," *Working Paper*, 2019.

Figure 1: Equilibrium in the marriage market

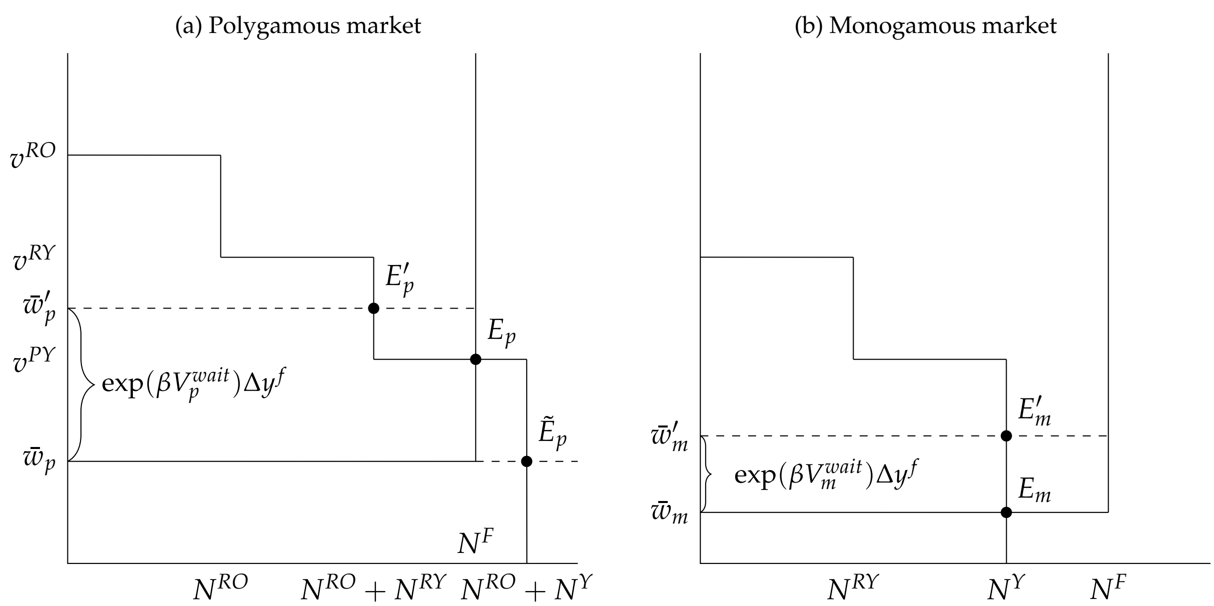

Figure 2: OLS correlation between Boko Haram and bride Gini

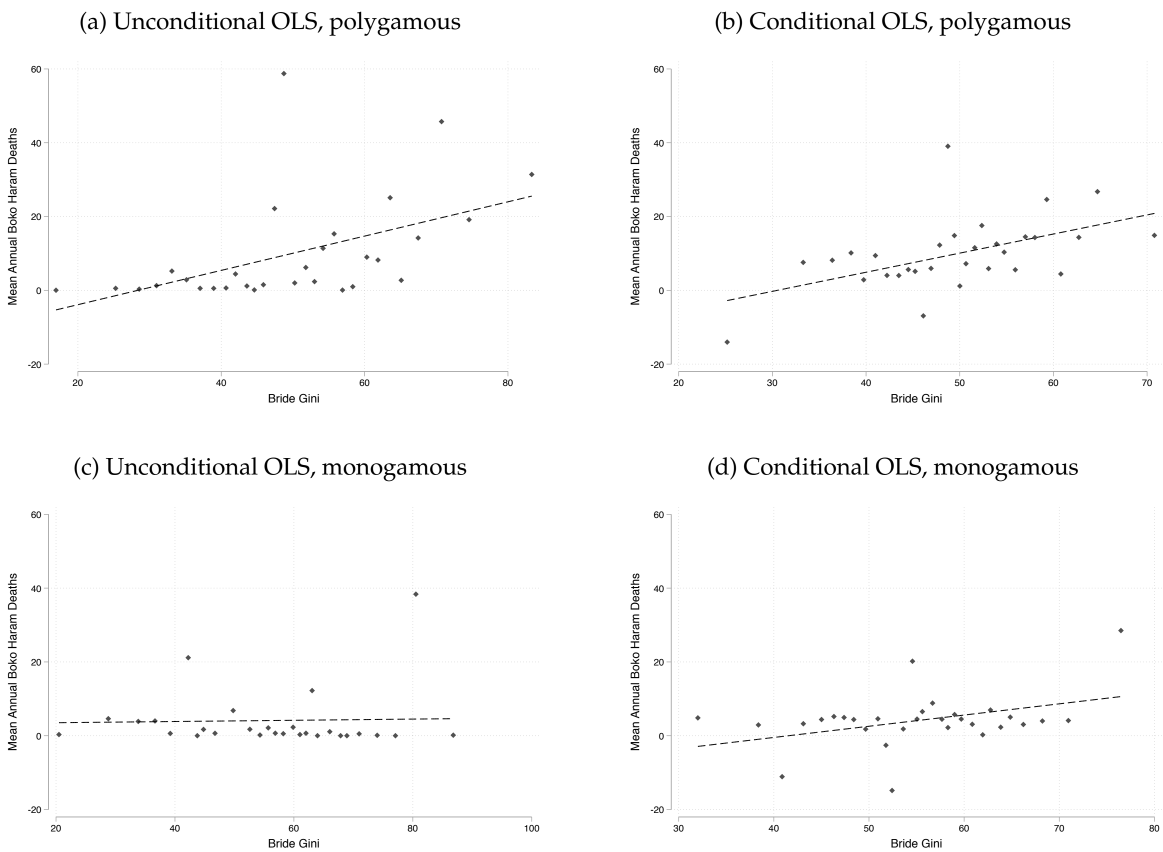

(a) Unconditional OLS, polygamous

(b) Conditional OLS, polygamous

(c) Unconditional OLS, monogamous

(d) Conditional OLS, monogamous

Note: Figure shows the residualized relationship (partial correlation) between the cluster-level mean annual Boko Haram-related fatalities and the bride Gini separately for monogamous and polygamous clusters. Scatterplots are binned using 30 quantiles of the distribution of the Bride Gini. Panel A shows the unconditional relationship while Panel B controls for DHS round fixed effects, ethnic homeland fixed effects, slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim.

Figure 3: OLS correlation between Boko Haram and adolescent rainfall by polygamy

(a) Polygamous areas

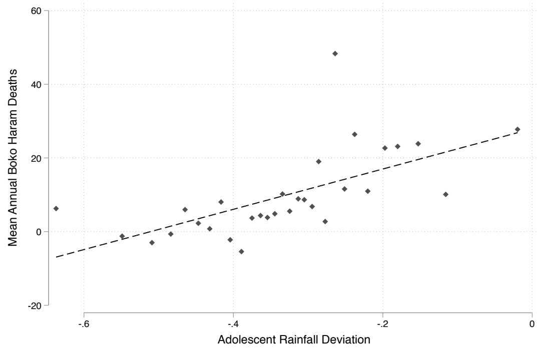

(b) Monogamous areas

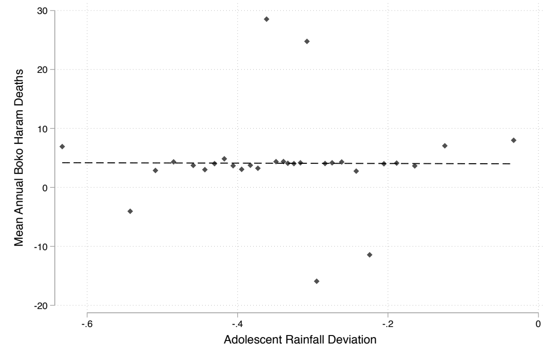

Note: Figure shows the residualized relationship (partial correlation) between the cluster-level mean annual Boko Haram-related fatalities and the mean annual rainfall for women between the ages of 12-16. Scatterplots are binned using 30 quantiles of the distribution of the bride Gini. Panel A estimates the relationship for the subsample of polygamous clusters while Panel B does the same for monogamous clusters. All plots include DHS round fixed effects and controls for slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim.

Figure 4: OLS correlation between bride Gini and adolescent rainfall by polygamy

(a) Polygamous areas

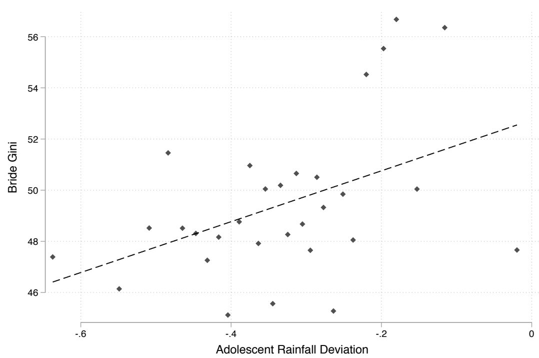

(b) Monogamous areas

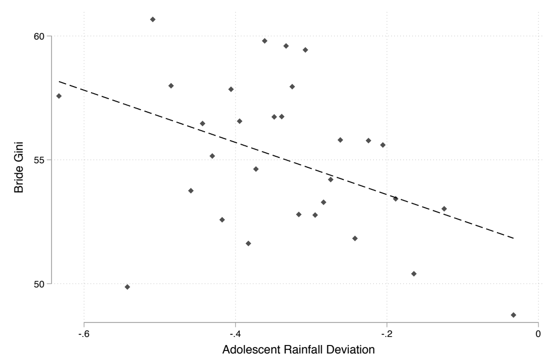

Note: Figure shows the residualized relationship (partial correlation) between the cluster-level bride Gini and the mean annual rainfall for women when they were ages 12-16. Scatterplots are binned using 30 quantiles of the distribution of the bride Gini. Panel A estimates the relationship for the subsample of polygamous clusters while Panel B does the same for monogamous clusters. All plots include DHS round fixed effects and controls for slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim.

Figure 5: Marriage market with age preferences

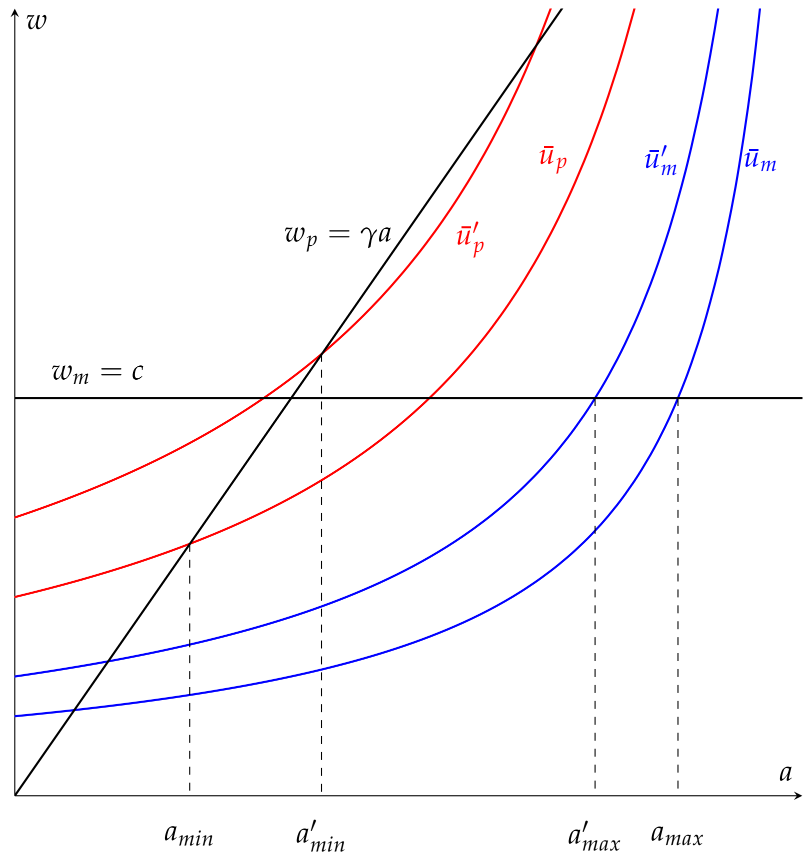

Figure 6: Placebo test of pre- and post-adolescent rainfall shocks

(a) Distribution of coefficients

(b) Time-path of coefficients

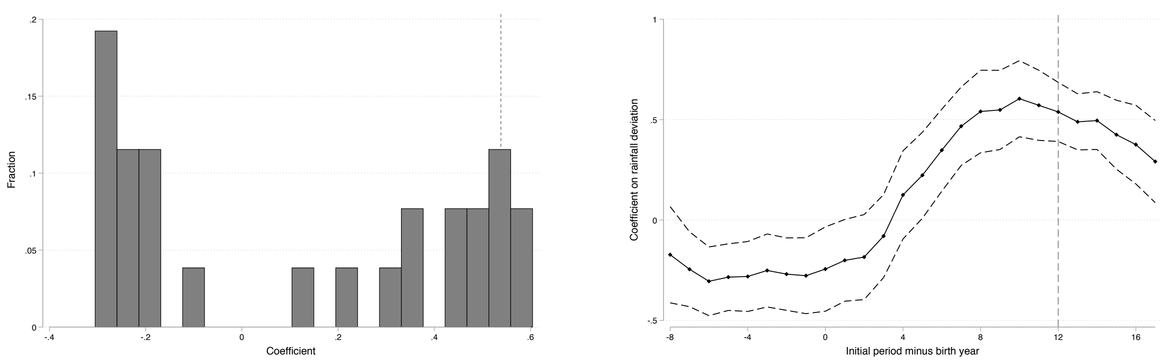

Note: Figure plots the coefficients from individual-level regressions of age of marriage on adolescent rainfall shocks s^{τ}_{ijc}, where τ refers to the first year, relative to the year of birth, of a four year window over which the shock is calculated. All regressions include DHS round fixed effects, grid-cell fixed effects, and controls for current rainfall deviation, a quadratic in age, village slope, and indicators for Muslim and Hausa. Coefficients on s^{τ}_{ijc} are plotted for each τ. Panel A presents the distribution of these coefficients, while Panel B plots their time-path, with a vertical line at age 16. 95% confidence intervals are constructed using standard errors clustered at the grid-cell level.

Figure 7: Adolescent rainfall and marriage hazard by age and polygamy status

(a) Polygamous areas

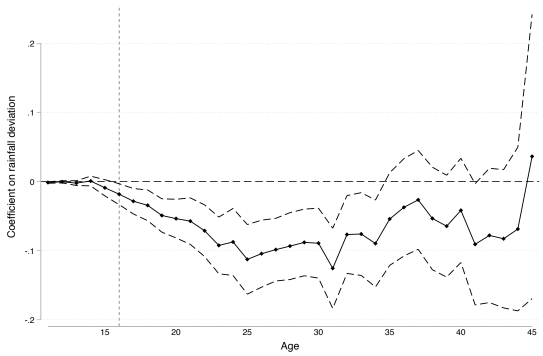

(b) Monogamous areas

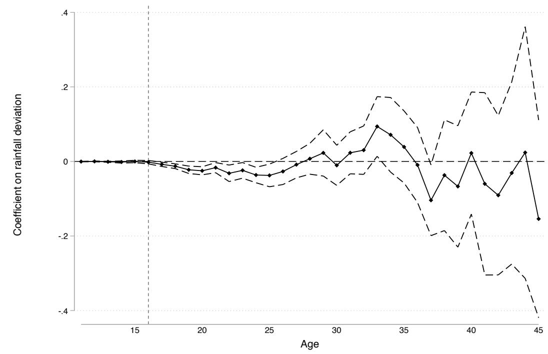

Note: Figure shows coefficients from an individual-level hazard regression of a marriage dummy on mean annual rainfall deviation from ages 12-16 interacted with age dummies, controlling for grid-cell and cohort fixed effects in a sample of women aged 10-45. 95% confidence bands, shown in dashed lines, are calculated using standard errors clustered at the grid-cell level. Panel A conducts the estimation on the subsample of polygamous clusters while Panel B does the same for monogamous clusters. Vertical line indicates age 16.

Figure 8: Placebo test for marriage expenditures

(a) Distribution of coefficients

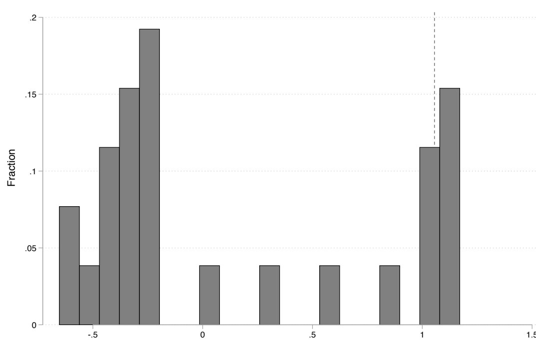

(b) Time-path of coefficients

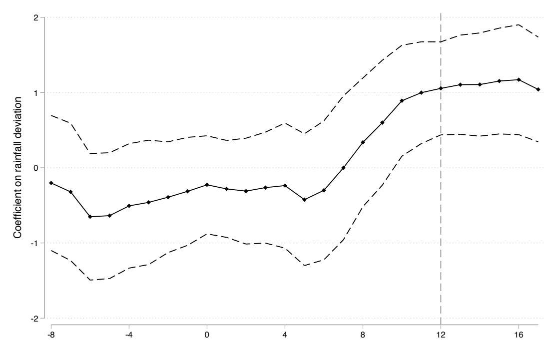

Note Figure plots the time-path of coefficients from regressions of log marriage expenditures on village-level mean adolescent rainfall shocks $s^{\tau}_{j,r}$, controlling for current rainfall and average female age. τ refers to the first year, relative to the year of birth, of a four year window over which the shock is calculated. All regressions are estimated on the subsample of polygamous villages. 95% confidence intervals are constructed using standard errors clustered at the grid-cell level.

Table 1: Summary statistics

Year	2008			2013		
Cluster	Monogamous	Polygamous	All	Monogamous	Polygamous	All
Cluster-level, $N = 1,768$						
Bride Gini	0.52	0.48	0.50	0.57	0.51	0.54
	(0.15)	(0.15)	(0.15)	(0.15)	(0.15)	(0.15)
Polygamous	0.00	1.00	0.64	0.00	1.00	0.59
	(0.00)	(0.00)	(0.48)	(0.00)	(0.00)	(0.49)
Boko Haram active	0.18	0.19	0.18	0.27	0.25	0.26
	(0.38)	(0.39)	(0.39)	(0.44)	(0.44)	(0.44)
Boko Haram deaths	1.48	4.34	3.30	6.48	15.68	11.95
	(15.69)	(25.09)	(22.16)	(49.22)	(65.41)	(59.52)
Boko Haram kidnapping	0.03	0.04	0.04	0.03	0.03	0.03
	(0.17)	(0.20)	(0.19)	(0.18)	(0.17)	(0.17)
Wealth Gini	0.15	0.19	0.17	0.09	0.10	0.10
	(0.07)	(0.06)	(0.07)	(0.03)	(0.03)	(0.03)
Wealth index	2.14	1.37	1.65	3.04	2.17	2.52
	(0.81)	(0.82)	(0.90)	(0.75)	(0.80)	(0.89)
Population density in 2005	2689.12	967.67	1594.95	2123.65	1170.03	1556.54
	(5785.82)	(3142.80)	(4373.86)	(4382.79)	(4069.38)	(4222.86)
Male, $N = 32,497$						
Age of marriage	27.34	24.05	24.99	27.21	24.20	25.15
	(6.14)	(5.43)	(5.84)	(5.72)	(5.19)	(5.54)
Unmarried	0.53	0.40	0.44	0.58	0.45	0.50
	(0.50)	(0.49)	(0.50)	(0.49)	(0.50)	(0.50)
One wife	0.47	0.45	0.45	0.42	0.42	0.42
	(0.50)	(0.50)	(0.50)	(0.49)	(0.49)	(0.49)
Two wives	0.00	0.13	0.09	0.00	0.12	0.08
	(0.00)	(0.34)	(0.29)	(0.00)	(0.33)	(0.26)
Three wives	0.00	0.02	0.01	0.00	0.01	0.01
	(0.00)	(0.13)	(0.11)	(0.00)	(0.11)	(0.09)
Four wives	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.06)	(0.05)	(0.00)	(0.05)	(0.04)
Muslim	0.18	0.59	0.45	0.19	0.70	0.51
	(0.39)	(0.49)	(0.50)	(0.40)	(0.46)	(0.50)
Hausa	0.05	0.32	0.23	0.04	0.41	0.27
	(0.22)	(0.47)	(0.42)	(0.21)	(0.49)	(0.45)
Kanuri	0.00	0.03	0.02	0.01	0.02	0.02
	(0.07)	(0.17)	(0.14)	(0.10)	(0.14)	(0.13)
Female, $N = 72,009$						
Age of marriage	20.14	16.92	17.81	20.23	16.70	17.77
	(5.16)	(4.26)	(4.75)	(5.09)	(4.07)	(4.69)
Polygamous	0.14	0.40	0.33	0.13	0.41	0.33
	(0.35)	(0.49)	(0.47)	(0.34)	(0.49)	(0.47)
Number of co-wives	1.19	1.51	1.42	1.17	1.51	1.41
	(0.59)	(0.74)	(0.72)	(0.53)	(0.72)	(0.69)
Muslim	0.17	0.58	0.44	0.19	0.70	0.52
	(0.38)	(0.49)	(0.50)	(0.40)	(0.46)	(0.50)
Hausa	0.04	0.31	0.22	0.04	0.41	0.28
	(0.21)	(0.46)	(0.42)	(0.20)	(0.49)	(0.45)
Kanuri	0.01	0.03	0.02	0.01	0.02	0.02
	(0.09)	(0.17)	(0.15)	(0.11)	(0.14)	(0.13)

Table displays means of variables with standard deviations in parentheses. Sample is a DHS repeated cross section at cluster, male, and female level, as indicated. Indicated sample sizes are pooled for all years and cluster types. Population density is measured as residents per square kilometer.

Table 2: The effect of marriage inequality on violence: baseline results

Dependent variable	Any Event (1)	Total Events (2)	Total Deaths (3)	$\log(1 + deaths)$ (4)	$\log(1 + events)$ (5)
Panel A: OLS					
Bride Gini	0.003***	0.042*	0.477*	0.014***	0.007***
	(0.001)	(0.025)	(0.270)	(0.005)	(0.003)
R^2	0.264	0.078	0.088	0.201	0.205
Panel B: Reduced form					
Adolescent rainfall deviation × Polygamous	0.275**	1.800**	20.086**	0.856***	0.407***
	(0.117)	(0.804)	(8.576)	(0.313)	(0.146)
R^2	0.278	0.067	0.078	0.190	0.193
Panel C: Two-stage least squares					
Bride Gini	0.021**	0.140**	1.565**	0.067***	0.032***
	(0.009)	(0.064)	(0.693)	(0.025)	(0.012)
Mean Dep. Var.	0.221	0.705	7.602	0.382	0.187
Dependent variable	Bride Gini				
Panel D: First-stage					
Adolescent rainfall deviation × Polygamous	12.833***				
	(2.689)				
Mean Dep. Var.	51.601				
R^2	0.524				
Kleibergen-Paap F-statistic	22.776				
Observations	1768	1768	1768	1768	1768

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. In Panels A, B, and C, outcome variable is (1) dummy variable indicating any Boko Haram related events, (2) the average annual number of Boko Haram events, or (3) average annual number of Boko Haram fatalities within 20 km of the DHS cluster. In Panel D, outcome variable is the Gini coefficient of wives in the cluster. All regressions include DHS round fixed effects and a full set of controls. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: The effect of marriage inequality on violence: robustness to specification

Dependent variable	Boko Haram deaths				
	(1)	(2)	(3)	(4)	(5)
Panel A: OLS					
Bride Gini	0.250*	0.477*	0.372*	0.429**	0.355*
	(0.132)	(0.270)	(0.194)	(0.203)	(0.204)
R^2	0.007	0.088	0.361	0.389	0.441
Panel B: Reduced form					
Adolescent rainfall deviation × Polygamous	12.761**	20.086**	9.760**	19.126***	8.059**
	(6.270)	(8.576)	(4.315)	(5.710)	(4.031)
R^2	0.008	0.078	0.354	0.386	0.436
Panel C: Two-stage least squares					
Bride Gini	0.674**	1.565**	0.994**	1.620***	0.810*
	(0.339)	(0.693)	(0.483)	(0.539)	(0.426)
Mean Dep. Var.	7.602				
Dependent variable	Bride Gini				
Panel D: First-stage					
Adolescent rainfall deviation × Polygamous	18.941***	12.833***	9.821***	11.809***	9.945***
	(3.149)	(2.689)	(2.891)	(2.908)	(3.101)
Mean Dep. Var	51.601				
R^2	0.051	0.524	0.555	0.573	0.587
Kleibergen-Paap F-statistic	36.175	22.776	11.542	16.489	10.284
Observations	1768	1768	1768	1768	1768
Controls	No	Yes	Yes	Yes	Yes
DHS Round FE	No	Yes	Yes	Yes	Yes
Tribe FE	No	No	No	Yes	Yes
State FE	No	No	Yes	No	Yes

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. Outcome variable is either the average annual number of Boko Haram related fatalities within 20 km of the DHS cluster (Panels A, B, C) or the Gini coefficient of wives in the cluster (Panel D). Tribe FE indicate ethnic homeland fixed effects. State FE indicate Nigerian state fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: The effect of adolescent rainfall shocks on spousal age gaps

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Polygamous			Monogamous		
Adolescent rainfall deviation	-0.645	0.013	-0.395	-0.863	-1.959**	-2.724**
	(0.481)	(0.639)	(0.814)	(0.585)	(0.951)	(1.100)
Mean Dep. Var	10.853			9.181		
Observations	1120	1120	1120	648	648	648
R^2	0.375	0.467	0.532	0.260	0.379	0.433
Controls	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	No	No	Yes	No	No	Yes
State FE	No	Yes	Yes	No	Yes	Yes

Standard errors are clustered at the grid-cell level. Outcome variable average spousal age gap the DHS cluster, as indicated. Adolescent rainfall deviation is the cluster-level mean of individual annual rainfall deviations from the long-run average between the ages of 12-16. Tribe FE indicate ethnic homeland fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: The effect of adolescent rainfall on marriage supply

Outcome	Age of marriage				Child marriage			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Panel A: OLS				
Adolescent rainfall	1.029***	1.325***	0.586***	0.542***	-0.071***	-0.070***	-0.045*	-0.039
	(0.107)	(0.131)	(0.072)	(0.075)	(0.025)	(0.027)	(0.026)	(0.027)
Mean Dep. Var.	17.679				0.217			
Observations	59652	59652	59652	59652	9874	9874	9874	9874
R^2	0.223	0.230	0.301	0.312	0.269	0.272	0.382	0.493
				Panel B: Heckman selection				
Adolescent rainfall	1.272***	1.628***	0.885***	0.831***				
	(0.102)	(0.116)	(0.084)	(0.091)				
λ	3.006***	3.057***	1.891***	1.675***				
	(0.290)	(0.278)	(0.154)	(0.171)				
Observations	59652	59652	59652	59652				
R^2	0.242	0.249	0.308	0.317				
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Grid-cell FE	No	No	No	Yes	No	No	No	Yes
State FE	No	No	Yes	No	No	No	Yes	No

Standard errors are clustered at the grid-cell level. Outcome variable is either age of marriage, or a child marriage indicator, given in table header. Independent variable is average annual rainfall deviation in the years between age 12 and 16. Sample is a repeated cross-section of individual women between ages 15 and 49. Fixed effects are for DHS round, cohort, state, and/or grid-cell, as indicated. Controls are current rainfall deviation, a quadratic in age, and dummies for Muslim and Hausa. Age of marriage includes the entire sample of ever-married women. For child marriage, the sample includes only women below age 18 and the outcome variable is a dummy marriage indicator. Panel A gives results for OLS, while Panel B gives results for two-step Heckman correction. The first step predicts ever-married with the selection variables age, age-squared, education level, year of birth, rural, muslim, Hausa, and Kanuri. λ is the inverse Mills ratio. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Match type and adolescent rainfall by polygamy

Outcome	Husband's wealth				Polygamous at marriage			
Cluster type	Polygamous		Monogamous		Polygamous		Monogamous	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Adolescent rainfall	0.077***	0.032**	0.010	0.009	0.011	0.015*	-0.010	-0.001
	(0.024)	(0.014)	(0.034)	(0.018)	(0.008)	(0.008)	(0.009)	(0.009)
Mean Dep. Var.	2.137		3.100		0.234		0.097	
Observations	40442	40442	15589	15589	40442	40442	15589	15589
R^2	0.340	0.485	0.359	0.516	0.020	0.033	0.050	0.068
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell FE	No	Yes	No	Yes	No	Yes	No	Yes
State FE	Yes	No	Yes	No	Yes	No	Yes	No

Standard errors are clustered at the grid-cell level. Outcome variable is either household wealth after marriage or an indicator for a married husband at the time of marriage. Estimation sample is all married women ages 15-49, and is split between monogamous and polygamous villages. Adolescent rainfall deviation is the mean annual rainfall deviation from the long-run average between the ages of 12-16. Regressions include either state or grid-cell fixed effects, as indicated. Controls are current rainfall deviation, a quadratic in age, and dummies for Muslim and Hausa. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 7: The effect of adolescent rainfall on marriage expenditures

Sample	Full			Polygamous			Monogamous		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Rainfall deviation	0.688*	0.676**	0.199	0.915***	0.709**	0.818*	0.528	0.642	-0.076
	(0.391)	(0.311)	(0.373)	(0.336)	(0.355)	(0.480)	(0.661)	(0.471)	(0.586)
Observations	1132	1132	1132	723	671	671	409	409	409
R^2	0.014	0.108	0.257	0.023	0.135	0.345	0.013	0.101	0.346
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
State FE	No	No	Yes	No	No	Yes	No	No	Yes

Standard errors are clustered at the grid-cell level. Outcome variable is the log of marriage expenditures at the household level. Independent variable is mean adolescent rainfall shock in the GHS cluster. Sample is pooled across all three GHS waves (2010, 2013, 2015) and includes every household incurring a wedding expenditure. All models include year dummies. Controls are rainfall deviation in 2010 and 2009, share muslim, distance to roads, distance to national borders, and average age of women. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: Adolescent rainfall, abductions, and violence against women

Outcome	Abductions			Violence against women		
	(1)	(2)	(3)	(4)	(5)	(6)
Adolescent rainfall deviation × Polygamous	0.060**	0.091***	0.066**	0.038	0.068**	0.043
	(0.029)	(0.033)	(0.030)	(0.024)	(0.029)	(0.027)
Mean Dep. Var.	0.033			0.041		
Observations	1768	1768	1768	1768	1768	1768
R^2	0.300	0.336	0.380	0.436	0.488	0.522
Controls	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	No	Yes	Yes	No	Yes	Yes
State FE	Yes	No	Yes	Yes	No	Yes

Standard errors are clustered at the grid-cell level. Outcome variable is a dummy variable indicating that the cluster experienced any Boko Haram abduction or attack against women in the period, as indicated. Adolescent rainfall deviation is the cluster-level mean of individual annual rainfall deviations from the long-run average between the ages of 12-16. Tribe FE indicate ethnic homeland fixed effects. Controls are slope, standardized rainfall deviation in the survey year, lagged rainfall deviation, population density in 2005, the wealth index Gini, distance to national borders, average age of men and women, share muslim, and share Hausa. All models also include the total number of Boko Haram attacks on the right hand side. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A Appendix figures

Figure A1: Polygamy rate

(a) 2008

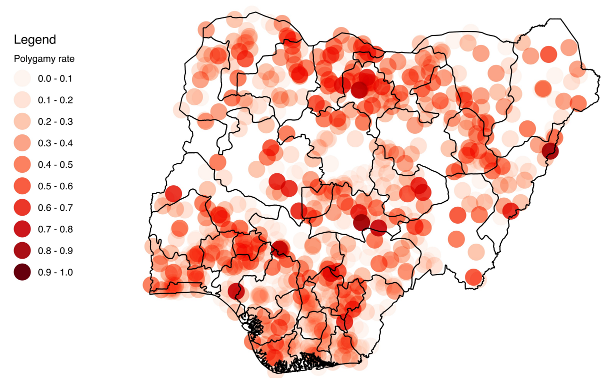

(b) 2013

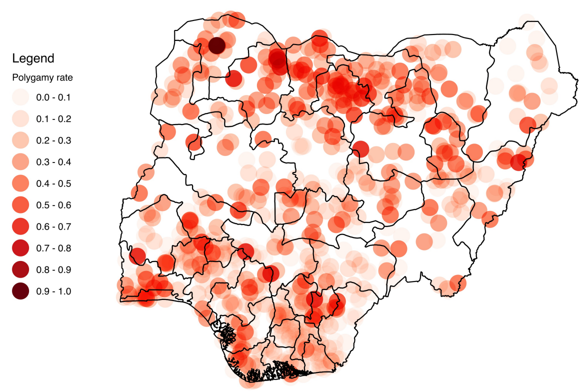

Note: Figure shows the geographic distribution of the polygamy rate among men in each DHS cluster, for DHS rounds 2008 and 2013.

Figure A2: Boko Haram deaths

(a) 2008-2012

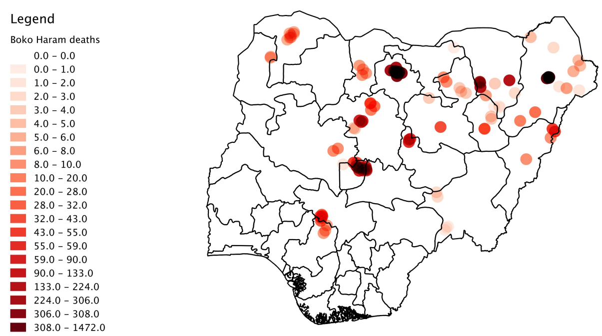

(b) 2013-2016

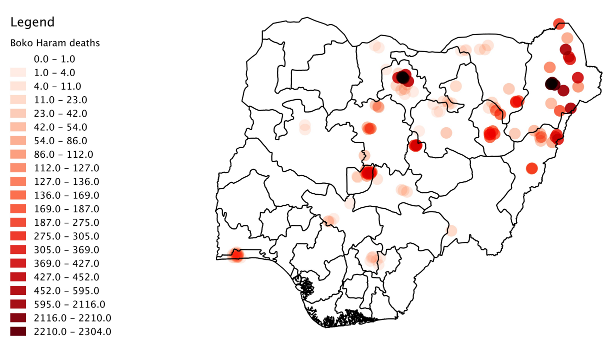

Note: Figure shows the geographic distribution of average annual Boko Haram-related fatalities in each DHS cluster, over the periods 2008-2012 and 2013-2016, corresponding to DHS rounds 2008 and 2013.

Figure A3: Distributions of key cluster-level variables

(a) Marriage inequality

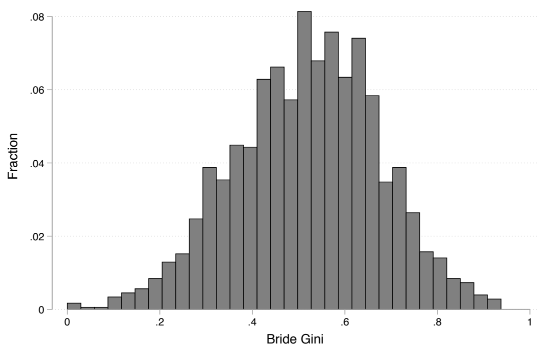

(b) Adolescent rainfall shock

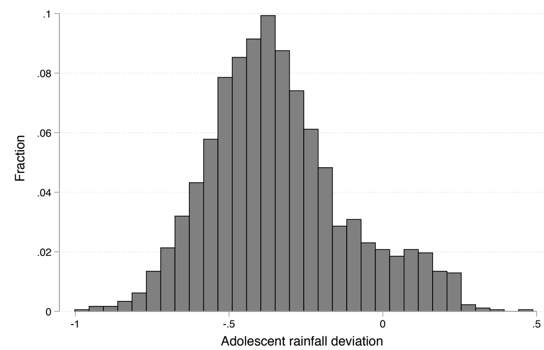

(c) Polygamy rate

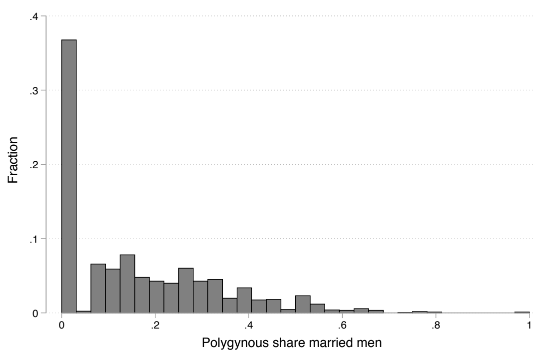

(d) Number of wives

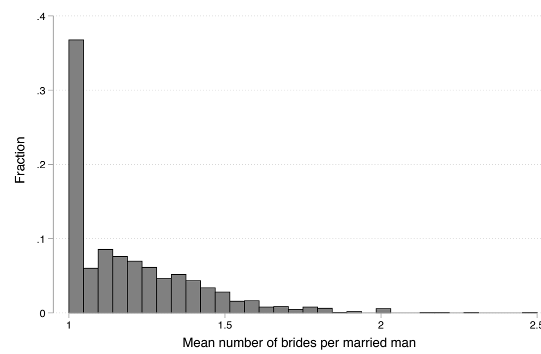

Note: Figure plots the frequency distribution of the following cluster-level variables: a) Bride Gini coefficient, b) cluster mean of individual rainfall deviations in the adolescent period of women, c) the share of married men in polygamous households, and d) the mean number of brides per married man.

Figure A4: OLS correlation between bride Gini and polygamy

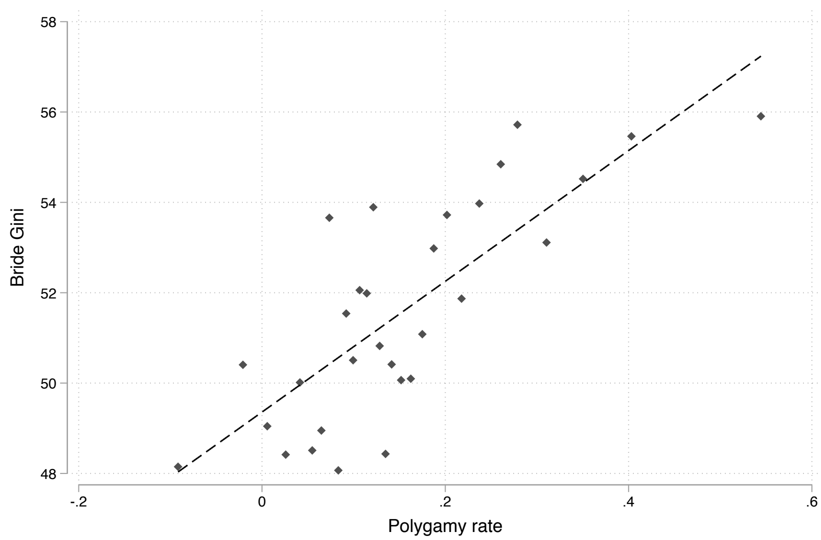

Note: Figure shows the residualized relationship (partial correlation) between the cluster-level bride Gini and the polygamy rate among men. Scatterplots are binned using 30 quantiles of the distribution of the polygamy rate. Panel A controls for DHS round fixed effects, ethnic homeland fixed effects, slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim.

Figure A5: Market share by polygamy and age

(a) Monogamous areas, under 35

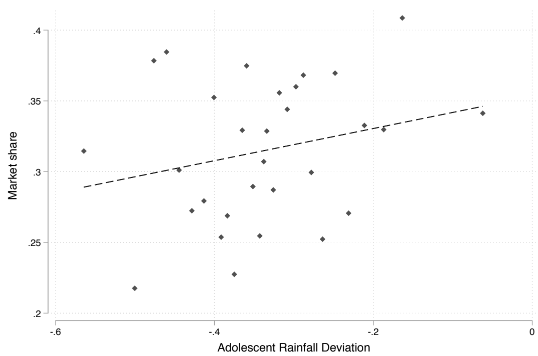

(b) Monogamous areas, over 35

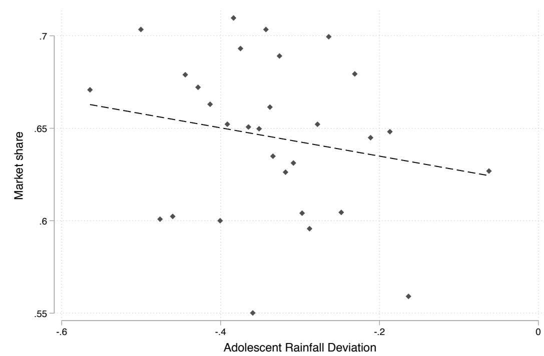

(c) Polygamous areas, under 35

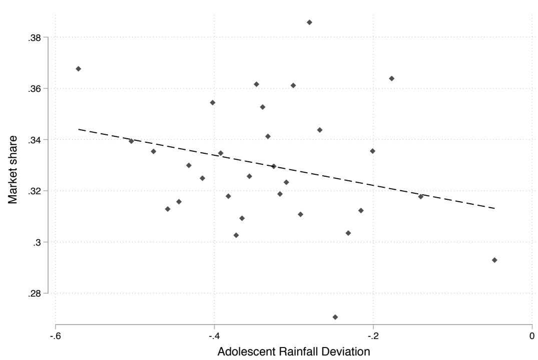

(d) Polygamous areas, over 35

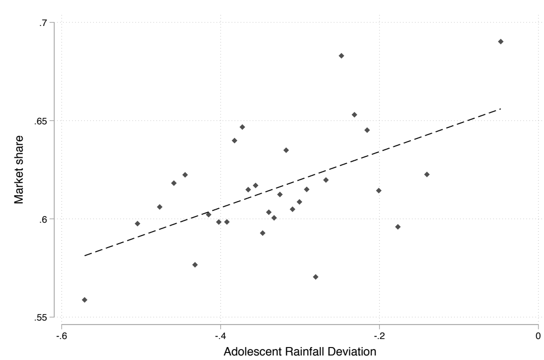

Note: Figure shows the residualized relationship (partial correlation) between the marriage market share of over and under 35 year old men and the mean annual rainfall of women between the ages of 12-16. Scatterplots are binned using 30 quantiles of the distribution of the bride Gini. Panels A and B estimates the relationship for the subsample of monogamous clusters, while Panels C and D do the same for polygamous clusters. All plots include DHS round and ethnic homeland fixed effects, and the full set of control variables.

Figure A6: Spousal age gaps and adolescent rainfall shocks

(a) Polygamous areas (b) Monogamous areas

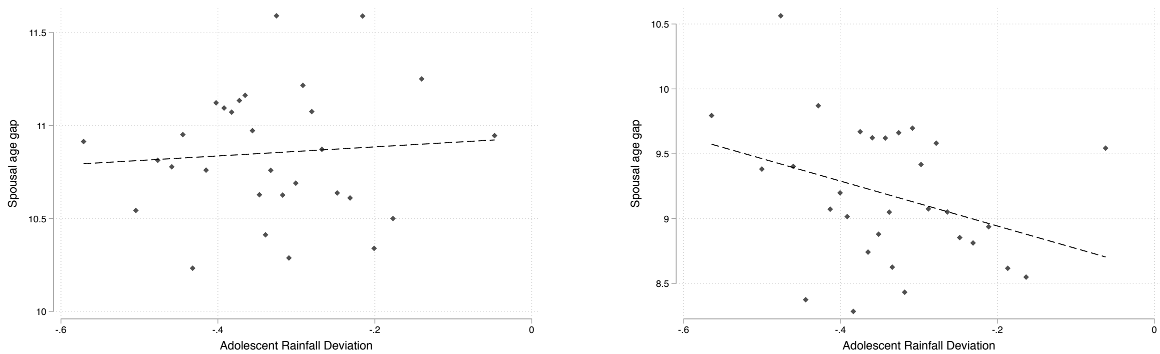

Note: Figure shows the residualized relationship (partial correlation) between the average spousal age gap and the mean annual rainfall of women between the ages of 12-16. Scatterplots are binned using 30 quantiles of the distribution of the bride Gini. Panel A estimates the relationship for the subsample of polygamous clusters while Panel B does the same for monogamous clusters. All plots include DHS round and ethnic homeland fixed effects, and the full set of control variables.

Figure A7: Market share and Boko Haram activity

(a) Under 35 (b) Over 35

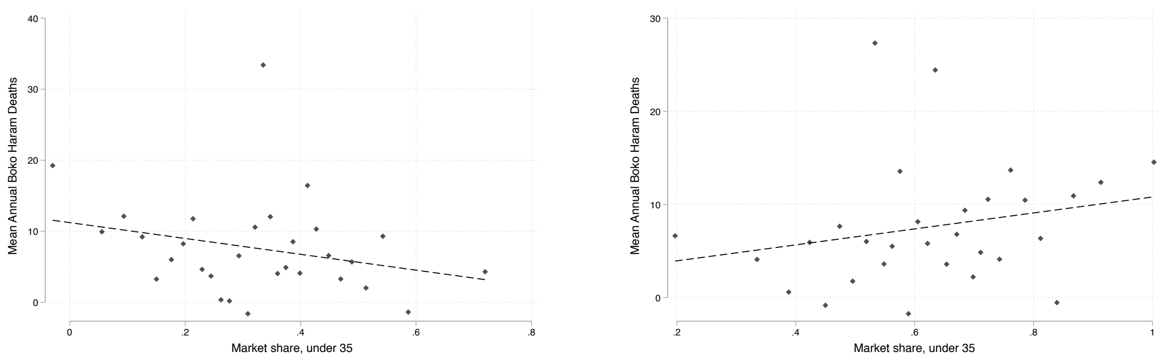

Note: Figure shows the residualized relationship (partial correlation) between mean annual Boko Haram fatalities and the marriage market share of over and under 35 year old men. Scatterplots are binned using 30 quantiles of the distribution of the bride Gini. Panel A estimates the relationship using market share of under 35 men, while Panel B does the same for using the market share of over 35 men. All plots include DHS round and ethnic homeland fixed effects, and the full set of control variables.

Figure A8: Placebo test: marriage inequality

(a) Distribution of coefficients (b) Time-path of coefficients

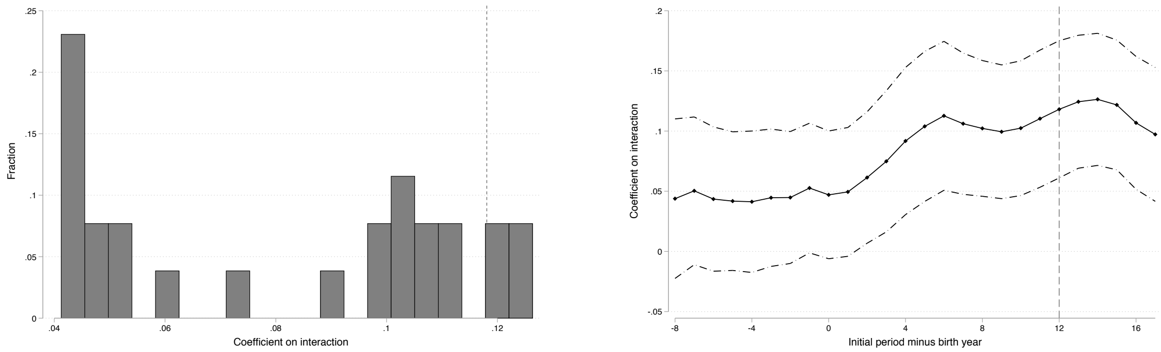

Note: Figure plots the coefficients from cluster-level regressions of Bride Gini on village-level mean adolescent rainfall shocks s_{jr}^{τ}, where τ refers to the first year, relative to the year of birth, of a four year window over which the shock is calculated. All regressions include DHS round fixed effects, ethnic homeland fixed effects, and controls for slope, standardized rainfall deviation in the survey year, lagged rainfall deviation, population density in 2005, the wealth index Gini, distance to national borders, average age of women, share muslim, and share Hausa. Coefficients on the interaction between s_{jr}^{τ} and p_{jr} are plotted for each τ. Panel A presents the distribution of these coefficients, while Panel B plots their time-path, with a vertical line at age 16. 95% confidence intervals are constructed using standard errors clustered at the grid-cell level.

Figure A9: Placebo test: Boko Haram

(a) Distribution of coefficients (b) Time-path of coefficients

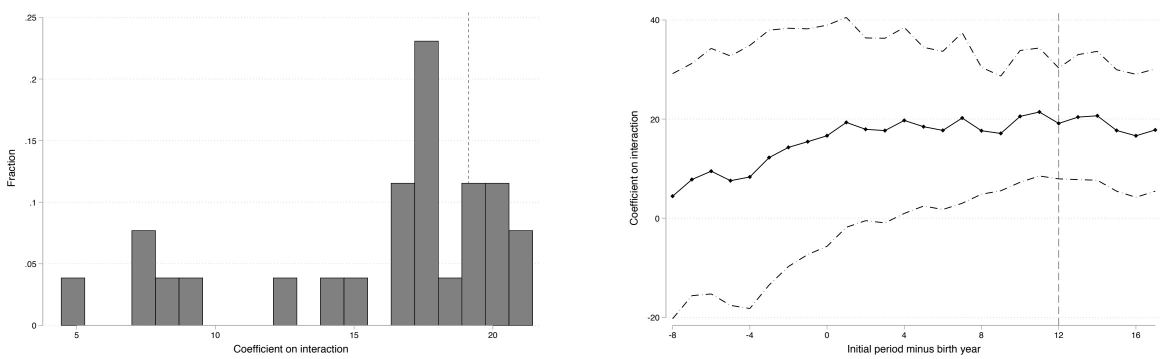

Note: Figure plots the coefficients from cluster-level regressions of mean annual Boko Haram fatalities on village-level mean adolescent rainfall shocks s_{jr}^{τ}, where τ refers to the first year, relative to the year of birth, of a four year window over which the shock is calculated. All regressions include DHS round fixed effects, ethnic homeland fixed effects, and controls for slope, standardized rainfall deviation in the survey year, lagged rainfall deviation, population density in 2005, the wealth index Gini, distance to national borders, average age of women, share muslim, and share Hausa. Coefficients on the interaction between s_{jr}^{τ} and p_{jr} are plotted for each τ. Panel A presents the distribution of these coefficients, while Panel B plots their time-path, with a vertical line at age 16. 95% confidence intervals are constructed using standard errors clustered at the grid-cell level.

B Appendix Tables

Table B1: Balance tests

Dependent variable	s_j (1)	$s_j \times p_j$ (2)	r_j (3)	$r_j \times p_j$ (4)	Observations (5)	R^2 (6)
Income						
Wealth index	0.486**	-0.036	0.010	0.093	1768	0.650
	(0.203)	(0.161)	(0.118)	(0.084)		
Wealth index Gini	-3.583**	2.881**	0.378	-0.271	1768	0.573
	(1.793)	(1.171)	(0.848)	(0.577)		
Demographics						
Male average age	0.565	-1.196	-0.022	0.831**	1768	0.295
	(0.851)	(0.809)	(0.448)	(0.364)		
Female average age	2.341***	-1.678***	-0.008	0.439*	1768	0.213
	(0.628)	(0.562)	(0.312)	(0.240)		
Muslim	0.022	-0.007	-0.043	0.052*	1768	0.855
	(0.042)	(0.057)	(0.032)	(0.031)		
Hausa	0.059*	-0.011	-0.016	-0.001	1768	0.809
	(0.030)	(0.035)	(0.024)	(0.020)		
Kanuri	-0.035***	-0.019	0	0.006	1768	0.591
	(0.013)	(0.016)	(0.012)	(0.010)		
Log of population density	0.875*	0.085	0.493**	-0.228	1768	0.643
	(0.510)	(0.290)	(0.241)	(0.164)		
Spatial characteristics						
Long-run mean annual rainfall	-15.729	71.795*	82.447**	-41.895**	1768	0.957
	(36.372)	(38.986)	(33.642)	(19.897)		
Long-run SD of annual rainfall	-2.959	8.997	24.083***	-10.075***	1768	0.958
	(7.700)	(5.504)	(4.558)	(3.413)		
Mean monthly temperature	0.190	-0.008	0.428***	-0.139***	1768	0.819
	(0.135)	(0.121)	(0.089)	(0.044)		
Altitude	-59.255*	-1.953	-67.551**	35.365***	1763	0.729
	(30.808)	(28.324)	(27.457)	(13.476)		
Distance to water body	-22.360**	-4.871	11.417	1.759	1768	0.948
	(8.951)	(7.973)	(7.142)	(3.964)		
Distance to border	10.667	5.317	-45.911***	6.114**	1768	0.915
	(7.311)	(6.985)	(6.146)	(2.949)		
Slope	0.045	-0.056	0.022	0.037	1768	0.544
	(0.094)	(0.111)	(0.060)	(0.036)		
Latitude	-0.260*	-0.070	-0.060	0.111**	1763	0.988
	(0.140)	(0.088)	(0.092)	(0.054)		
Conflict variables						
Riot deaths	-1.351	2.339*	1.388	-0.949	1768	0.235
	(1.100)	(1.332)	(1.052)	(0.738)		
Communal militia deaths	-0.789	0.676	1.051	-0.944	1768	0.295
	(1.669)	(1.803)	(1.667)	(1.051)		

Each cell represents the coefficients from a cluster-level regression of the dependent variable (specified in the first column) on the adolescent rainfall shock s_j, current rainfall deviation r_j, and their interactions with the polygamy dummy p_j. All specifications include DHS round dummies and state and ethnic homeland fixed effects. Standard errors in parentheses are clustered at the grid-cell level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B2: The effect of marriage inequality on violence: unmarried share as endogenous variable

Dependent variable	Boko Haram deaths				
	(1)	(2)	(3)	(4)	(5)
Panel A: OLS					
Share unmarried men	0.205*	0.534*	0.385**	0.440**	0.353*
	(0.114)	(0.308)	(0.193)	(0.200)	(0.197)
R^2	0.007	0.094	0.362	0.391	0.442
Panel B: Two-stage least squares					
Share unmarried men	0.585**	1.413**	0.975**	1.555***	0.860*
	(0.291)	(0.610)	(0.478)	(0.510)	(0.465)
Mean Dep. Var	7.602				
Dependent variable	Male unmarried rate				
Panel C: First-stage					
Adolescent rainfall deviation $\times$ Polygamous	21.801***	14.218***	10.012***	12.303***	9.365***
	(3.312)	(2.823)	(3.034)	(3.028)	(3.224)
Mean Dep. Var	46.149				
R^2	0.171	0.611	0.641	0.654	0.669
Kleibergen-Paap F-statistic	43.340	25.361	10.887	16.510	8.438
Observations	1768	1768	1768	1768	1768
Controls	No	Yes	Yes	Yes	Yes
DHS Round FE	No	Yes	Yes	Yes	Yes
Tribe FE	No	No	No	Yes	Yes
State FE	No	No	Yes	No	Yes

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. Outcome variable is either the average annual number of Boko Haram related fatalities within 20 km of the DHS cluster (Panels A, B) or the unmarried rate among men aged 15-59 (Panel C). Tribe FE indicate ethnic homeland fixed effects. State FE indicate Nigerian state fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B3: Reduced form and first stage: Robustness to current and lagged rainfall

Outcome	Bride Gini				Boko Haram deaths			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Adolescent rainfall deviation	-2.923	-1.550	-9.351***	-9.102***	-6.022	0.464	-1.598	14.820*
	(2.054)	(2.128)	(2.627)	(2.722)	(7.964)	(8.513)	(7.434)	(7.817)
Adolescent rainfall deviation × Polygamous	12.833***	10.985***	9.287***	11.369***	20.086**	10.937	10.160*	19.552**
	(2.689)	(2.873)	(3.082)	(3.275)	(8.576)	(8.040)	(5.707)	(8.412)
Current rainfall deviation	-3.190***	-4.846***	-0.770	-2.530	-3.395	-8.583***	0.256	-4.968
	(0.808)	(1.234)	(1.496)	(1.767)	(3.424)	(3.283)	(5.593)	(4.775)
Current rainfall deviation × Polygamous		2.389*	1.584	1.582		7.713*	0.004	0.415
		(1.350)	(1.313)	(1.488)		(3.957)	(3.058)	(2.421)
Lagged rainfall deviation	-2.803***	-2.214	0.747	0.179	12.592	12.377	7.081	3.161
	(0.735)	(1.365)	(1.383)	(1.463)	(8.669)	(10.694)	(6.484)	(4.703)
Lagged rainfall deviation × Polygamous		-0.580	-1.036	-1.163		1.265	-0.401	-0.844
		(1.561)	(1.548)	(1.665)		(3.824)	(4.621)	(6.144)
Observations	1768	1768	1768	1768	1768	1768	1768	1768
R^2	0.524	0.525	0.555	0.573	0.078	0.080	0.354	0.386
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	No	No	No	Yes	No	No	No	Yes
State FE	No	No	Yes	No	No	No	Yes	No

Standard errors are clustered at the grid-cell level. Outcome variable is either the mean annual Boko Haram-related fatalities, or the Bride Gini, as indicated in table header. Adolescent rainfall deviation is the cluster-level mean of individual annual rainfall deviations from the long-run average between the ages of 12-16. Tribe FE indicate ethnic homeland fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B4: Instrumental variables: robustness to current and lagged rainfall

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Bride Gini	1.089	1.179	1.143*	1.014*	1.779**	1.632**	0.946*	0.993
	(0.670)	(0.778)	(0.628)	(0.603)	(0.717)	(0.721)	(0.564)	(0.630)
Current rainfall deviation × Polygamous	5.991*		-1.472		-1.860		-1.353	
	(3.608)		(3.897)		(4.899)		(4.008)	
Lagged rainfall deviation × Polygamous		4.893		-0.193		-0.145		-1.839
		(4.286)		(4.572)		(6.330)		(5.704)
Kleibergen-Paap F-statistic	14.912	18.425	8.575	11.137	11.881	15.084	7.856	10.012
Anderson-Rubin p-value	0.085	0.123	0.027	0.070	0.003	0.016	0.069	0.109
Observations	1768	1768	1768	1768	1768	1768	1768	1768
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	No	No	Yes	Yes	No	No	Yes	Yes
Tribe FE	No	No	No	No	Yes	Yes	Yes	Yes

Standard errors are clustered at the grid-cell level. Outcome variable is the average annual number of Boko Haram-related fatalities within 20 km of the DHS cluster. All models are estimated with 2SLS, using the interaction between mean adolescent rainfall and polygamy status as an instrument for the Bride Gini. Tribe FE indicate ethnic homeland fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B5: Reduced form and first stage: Robustness to ethnicity and religion

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Boko Haram deaths								
Adolescent rainfall deviation × Polygamous	20.550**	19.756**	10.891*	8.573	18.354***	17.968***	11.881**	7.694**
	(8.451)	(8.088)	(5.890)	(5.626)	(5.893)	(5.892)	(5.080)	(3.430)
Adolescent rainfall deviation × Kanuri		2.411				1.200		
		(1.900)				(1.666)		
Adolescent rainfall deviation × Muslim			0.425**				0.325**	
			(0.196)				(0.145)	
Adolescent rainfall deviation × Hausa				0.936***				0.835***
				(0.289)				(0.263)
Observations	1768	1768	1768	1768	1768	1768	1768	1768
R^2	0.116	0.122	0.122	0.129	0.428	0.429	0.430	0.437
Panel B: Bride Gini								
Adolescent rainfall deviation × Polygamous	12.847***	12.988***	12.662***	11.161***	11.817***	11.962***	10.851***	9.219***
	(2.687)	(2.684)	(3.057)	(2.807)	(2.909)	(2.902)	(3.223)	(3.024)
Adolescent rainfall deviation × Kanuri		-0.429**				-0.450**		
		(0.182)				(0.227)		
Adolescent rainfall deviation × Muslim			0.008				0.049	
			(0.038)				(0.044)	
Adolescent rainfall deviation × Hausa				0.132***				0.203***
				(0.048)				(0.053)
Observations	1768	1768	1768	1768	1768	1768	1768	1768
R^2	0.524	0.526	0.524	0.526	0.573	0.574	0.573	0.577
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	No	No	No	No	Yes	Yes	Yes	Yes

Standard errors are clustered at the grid-cell level. Outcome variable is given in the panel header. Tribe FE indicate ethnic homeland fixed effects. Adolescent rainfall deviation is the cluster-level mean of individual annual rainfall deviations from the long-run average between the ages of 12-16. Demographic variables are the share of the village coming from Kanuri, Muslim, or Hausa groups, respectively, measured from 0 to 100. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, share Muslim, and share Kanuri.*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B6: The effect of marriage inequality on violence: non-Boko Haram violence

Dependent variable	Attacks			Fatalities		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS						
Bride Gini	0.029**	0.014	0.013	-0.048	-0.039	-0.063
	(0.014)	(0.014)	(0.014)	(0.072)	(0.092)	(0.098)
R^2	0.409	0.533	0.563	0.062	0.239	0.293
Panel B: Reduced form						
Adolescent rainfall deviation × Polygamous	-1.011	-1.086	-1.093	-5.280	-1.332	-3.048
	(3.774)	(3.457)	(3.862)	(6.934)	(5.185)	(5.343)
R^2	0.415	0.536	0.565	0.065	0.239	0.294
Panel C: Two-stage least squares						
Bride Gini	-0.079	-0.111	-0.110	-0.411	-0.136	-0.306
	(0.296)	(0.354)	(0.378)	(0.565)	(0.531)	(0.543)
Dependent variable	Bride Gini					
Panel D: First-stage						
Adolescent rainfall deviation × Polygamous	12.833***	9.821***	9.945***	12.833***	9.821***	9.945***
	(2.689)	(2.891)	(3.101)	(2.689)	(2.891)	(3.101)
R^2	0.524	0.555	0.587	0.524	0.555	0.587
Kleibergen-Paap F-statistic	22.776	11.542	10.284	22.776	11.542	10.284
Observations	1768	1768	1768	1768	1768	1768
Controls	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	No	No	Yes	No	No	Yes
State FE	No	Yes	Yes	No	Yes	Yes

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. In Panels A, B, and C, outcome variable is the average annual number of non-Boko Haram attacks (columns 1-3) or fatalities (columns 4-6) within 20 km of the DHS cluster. In Panel D, outcome variable is the Gini coefficient of wives in the cluster. All regressions include DHS round fixed effects and a full set of controls. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C Additional results

C.1 Inequality

Several auxiliary features of the model can also be tested directly in the data. Firstly, the woman's reservation policy implies a direct mapping between the ex-ante distribution of income, and the ex-post distribution of brides. Therefore, at the cluster level wealth inequality should be positively associated with marriage inequality, holding mean wealth fixed.[43] In Table C1 I consider the relationship between cluster-level wealth inequality and marriage inequality, estimating a regression of the bride Gini G^b_{jr} on the wealth Gini G^w_{jr}, and the standard set of controls and fixed effects used throughout the paper. Omitting the mean of the income distribution in (1) and (2), the correlation between income and bride inequality is negative. However, the sign flips once we include the mean level of wealth in column (3). The estimates from (3)-(7) imply that a one-point increase in G^w_{jr} is associated with a 0.15-0.5 point increase in G^b_{jr}. The unconditional specifications are negatively biased because richer communities have both less wealth inequality and more marriage inequality, on average. Once we condition on mean wealth, the prediction holds.

C.2 Wealth and brides

The model assumes that rich men who re-enter the market have higher willingness to pay, implying a positive correlation between wealth and brides. Table C2 and Figure C1 test the relationship between wealth and brides at the individual level. I regress the number of wives reported by a man on his wealth index and demographics, which include age, level of education, and fixed effects for self-reported ethnicity and religion. When I include only state fixed effects, the results suggest a negative correlation between male wealth and the number of brides (columns (1)-(4) of Table C2 and Panel A of Figure C1). However, these correlations are likely to be biased by unobserved local heterogeneity: less educated, poorer, more traditional areas are more likely to practice polygamy. I account for this heterogeneity with cluster fixed effects. Once we add these fixed effects in columns (5)-(8) and panel B of Figure C1, the sign flips in the expected direction. Accounting for market-specific heterogeneity, I find that a one point increase in the wealth index is associated with approximately 0.07 additional wives, significant at 1%. In columns (7) and (8) I also include the interaction of the wealth index with village-level polygamy status to establish that the wealth-brides gradient is larger polygamous areas. This is obvious, because in polygamous villages the upper bound on the number

[43] For intuition, consider the case where the income distribution is a mass point. Then all men are married, or not, and there is no marriage inequality. Marriage inequality can only exist if there is income inequality.

of women a rich man can accumulate is higher.

C.3 Age income profile

In Section 6.5, I extend the model to allow women to react to income shocks along the margins of both age and bride price. Here, predictions about the age profile of marriages are driven by the assumption that the age-wealth profile of marriageable men is positive under polygamy and flat under monogamy. Intuitively, under polygamy, older, richer men can re-enter the marriage market to compete with younger, poorer men. However, under monogamy, the old men who remain unmarried (and thus marriageable) are those too poor to afford a wife, flattening the age-income profile. Table C3 provides support for this assumption. Columns (1)-(3) and (4)-(6) demonstrate that, in both polygamous and monogamous markets, older men have significantly greater wealth. However, among the subsample of marriageable men – the unmarried – in monogamous markets, this correlation disappears.

Figure C1: Wealth and wives

(a) OLS (b) Cluster fixed effects

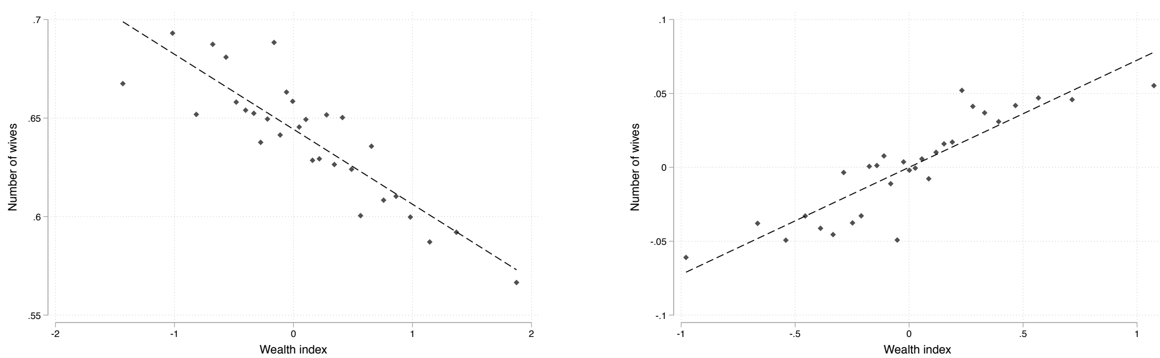

Note: Figure shows the residualized relationship (partial correlation) between the household-male level wealth score and number of wives. Scatterplots are binned using 30 quantiles of the wealth distribution. Panel A estimates the relationship using OLS, controlling only for education groups, age, and dummies for religion and ethnic group. Panel B includes the same controls, as well as DHS cluster fixed effects.

Figure C2: Wealth and age

(a) Polygamous

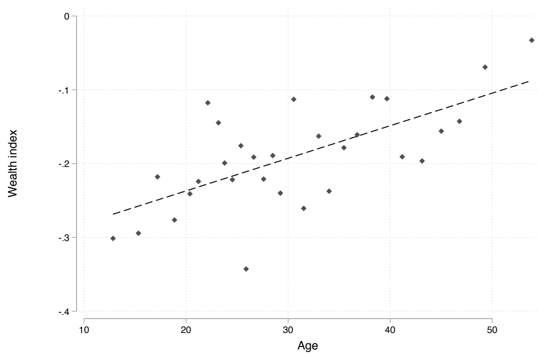

(b) Monogamous, married

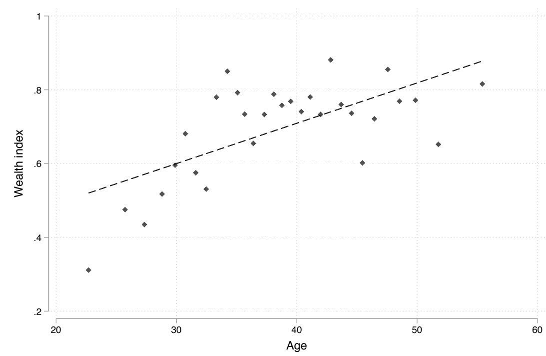

(c) Monogamous, unmarried

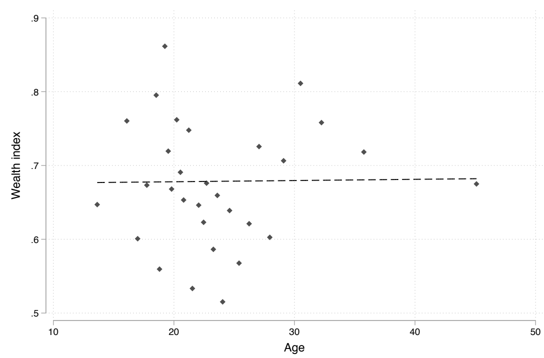

Note: Figure shows the residualized relationship (partial correlation) between the age and wealth score. Scatterplots are binned using 30 quantiles of the wealth distribution. All plots include dummies for education level and DHS round. Panel A estimates the relationship on the sample of polygamous clusters, Panel B for the subsample of married men in monogamous clusters, and Panel C for the subsample of unmarried men in monogamous clusters.

Table C1: Wealth inequality and marriage inequality

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Wealth index Gini	-0.213***	-0.158***	0.536***	0.234***	0.222**	0.173**	0.155**
	(0.080)	(0.061)	(0.091)	(0.066)	(0.104)	(0.070)	(0.074)
Mean Dep. Var.	51.601						
Observations	1768	1768	1768	1768	1768	1768	1768
R^2	0.022	0.511	0.154	0.556	0.262	0.579	0.605
Wealth control	No	No	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	No	No	No	No	No	No	Yes
State FE	No	No	No	No	Yes	Yes	Yes

Standard errors are clustered at the grid-cell level. Outcome variable is the Gini coefficient of brides among men age 15-59 in the DHS cluster. Independent variable is the Gini coefficient of household wealth scores within the village. Tribe FE indicate ethnic homeland fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, average age of men and women, distance to national borders, share Hausa, and share Muslim. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C2: Wealth and polygamy

Estimation	OLS				Cluster Fixed Effects			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Wealth Index	-0.112***	-0.068***	-0.038***	-0.027***	0.068***	0.073***	0.047***	0.050***
	(0.005)	(0.004)	(0.004)	(0.005)	(0.007)	(0.007)	(0.008)	(0.008)
Wealth Index × Polygamous							0.036***	0.037***
							(0.013)	(0.013)
Age		0.042***	0.042***	0.042***	0.042***	0.042***	0.042***	0.042***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Mean Dep. Var.	0.640							
Observations	34749	34749	34749	34749	34749	34749	34749	34749
R^2	0.028	0.489	0.511	0.515	0.561	0.567	0.561	0.567
Education controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Religion FE	No	No	Yes	Yes	No	Yes	No	Yes
Ethnicity FE	No	No	Yes	Yes	No	Yes	No	Yes
State FE	No	No	No	Yes	No	No	No	No
Cluster FE	No	No	No	No	Yes	Yes	Yes	Yes

Standard errors are clustered at the grid-cell level. Outcome variable is the number of wives in a sample of men aged 15-59. Ethnicity FE indicate self-reported ethnicity dummies, while religion FE indicates self-reported religious affiliation dummies. Education controls are included as dummies for 5 education levels, with no education as the omitted group. OLS specifications exclude cluster-level fixed effects, while cluster fixed effects specifications include them. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table C3: Wealth and age

Sample	Polygamous			Monogamous, all			Monogamous, unmarried		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Age	0.004***	0.002***	0.001***	0.006***	0.003***	0.002***	0.000	-0.002	-0.002
	(0.001)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
Mean Dep. Var.	2.326			3.201			3.194		
Observations	22866	22866	22866	11883	11883	11883	6691	6691	6691
R^2	0.353	0.545	0.823	0.267	0.536	0.786	0.214	0.504	0.791
Education controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Religion FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Ethnicity FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
State FE	No	Yes	No	No	Yes	No	No	Yes	No
Cluster FE	No	No	Yes	No	No	Yes	No	No	Yes

Standard errors are clustered at the cluster level. Outcome variable is the wealth index in a sample of men aged 15-59. Estimation sample in column (1)-(3) is all men in polygamous clusters, (4)-(6) is all men in monogamous clusters, and (7)-(9) is unmarried men in monogamous clusters. Ethnicity FE indicate self-reported ethnicity dummies, while religion FE indicates self-reported religious affiliation dummies. Education controls are included as dummies for 5 education levels, with no education as the omitted group. OLS specifications exclude cluster-level fixed effects, while cluster fixed effects specifications include them. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D Additional robustness tests

In this section, I conduct several additional robustness tests. First, I replicate the OLS, first stage, and reduced form results in a variety of subsamples. Secondly, I consider whether the results withstand varying the threshold used to define the polygamy dummy p_{jr}. Thirdly, I consider whether the results replicate across several additional conflict event data sources. Lastly, I account for measurement error in the adolescent rainfall shocks induced by female migration, replicating the primary individual and market-level marriage regressions with subsamples of women who arrived in their current village at early ages.

D.1 Robustness to sample definition

The effects should be more prominent in certain subsamples, and relatively less pronounced in others. In particular, the Boko Haram effect should be pronounced in Northern Nigeria, where the Islam is the dominant faith and the insurgency has taken hold. It should be relatively absent in the southern part of the country which lacks the necessary muslim population to support an extremist insurgency. We should also observe differential effects along the rural/urban axis, since marriage should be more responsive to adolescent rainfall conditions in primarily agricultural regions. Finally, we might also expect poorer villages to be more responsive, since the mechanisms in question might operate more strongly among the credit constrained.

In Table D1 I replicate the main market-level results, splitting the sample by north/south, rural/urban, and poor/rich.[44] Since the smaller samples substantially reduce statistical power, and may lead to an underpowered first stage, I omit the 2SLS results and provide only OLS, reduced form, and first stage. The results are generally as expected. Both first stage and reduced form effects are stronger in the North than in the South, where they are practically absent. The same is true in rural areas relative to urban ones. Lastly, poor areas drive the results relative to rich ones, which see null effects. However, the reduced form is only significant at 10% in the rural and poor subsamples, likely because of reduced power. Given this, the results should be interpreted with caution.

D.2 Robustness to polygamy threshold

I have argued throughout the paper that using a threshold of "any polygamous man" to define the existence of a cluster-level polygamy norm is preferable to using the equilibrium polygamy rate.

[44] Subnational region and rural/urban designation is defined by the Nigerian government in the DHS data. The wealth category of the village is defined as having a village-level mean wealth score below (poor) or above (rich) the median of the distribution of village-level mean wealth scores.

Nevertheless, I consider the robustness of the main results to alternate definitions of the polygamy threshold: 5%, 10%, 15% and 20% of men, as well as the raw polygamy rate. I also test robustness to using the woman-level data, which is drawn from a larger sample and thus more likely to capture instances where, by chance, the male sample excluded all polygamous men even in the presence of village polygamy. In Table D2 I re-estimate the first stage, reduced form, and 2SLS regressions using these thresholds for the polygamy rate. The results are not sensitive to this choice, and the estimated 2SLS coefficient changes only slightly.

D.3 Robustness to conflict data

Throughout the paper, I use conflict data from the Armed Conflict Location and Event Dataset (ACLED), a standard source in the empirical conflict literature. However, as noted by Donnay et al. (2018), it is preferable to combine information from numerous conflict event datasets rather than rely on a single source. I consider three additional sources of data which cover Africa's recent conflict history: the Global Terrorism Database (GTD), the Uppsala Conflict Data Program (UCDP), the Social Conflict Analysis Database (SCAD), all of which employ similar data-gathering methodologies. Table D3 replicates the main market-level specifications with the outcome variable as various different data sources, with the ACLED specification in column (1) for reference. The coefficients in the reduced form and 2SLS regressions are all of similar magnitude and remain significant at 5 or 1%.

D.4 Accounting for female migration

Throughout the paper, I assign each DHS cluster to it's nearest rainfall grid, and construct the individual-specific adolescent rainfall shocks using information on a woman's current location (DHS cluster) and year of birth. However, simply because a woman is surveyed in cluster j does not mean she was born and/or raised in that cluster. Women who migrated to the cluster in which they are surveyed from a different rainfall grid-cell will be assigned the wrong rainfall history. This introduces unavoidable measurement error, since the Nigerian DHS does not include place of origin. If early-life migration is correlated with outcomes of interest, this may introduce endogeneity issues beyond attenuation bias.

I address this problem by considering the subsample of women who migrated to the DHS cluster in which they were interviewed in time to experience the adolescent rainfall to which they are assigned. Unfortunately, data on age of migration is only available for DHS round 2008, so this reduces the sample substantially. I replicate both individual and market-level results for women who migrated by ages 12 and 16. Figure D1 provides the results of the market-level reduced form and first-stage estimation for the age 16 sample, while Figure D2 provides the results for the age 12 sample. The results generally echo those of corresponding full-sample estimation in Figures 3 and 4: positive slopes in polygamous subsample and zero or negative in the monogamous. Interestingly, it appears that, while the slopes for both the reduced for and first stage are negative in the monogamous subsamples when we use the full sample, they attenuate to zero when we use the age 12 and age 16 samples. This suggests that the negative effect of good rainfall on marriage inequality and Boko Haram activity in monogamous regions observed in the main results is actually being driven by movers.

The individual-level marriage hazard results are contained in Figures D3 and D4, which replicate the full-sample results of Figure 7 in the age 16 and age 12 samples, respectively. In general, the results are not much effected, and in fact the negative effect of the adolescent rainfall shock on marriage hazard in polygamous regions emerges even more prominently after age 16 than in the main results. This might be because of attenuation bias introduced by measurement error. Lastly, in Table D4 I re-estimate the age of marriage and child marriage specifications, using only the sample of women who lived in the cluster as of ages 12 or 16. The results appear to hold in this subsample, and are remarkably similar to the main effects. All told, these exercises suggest that the measurement error introduced by migration is not substantially affecting the main results.

D.5 Accounting for male migration

There is a concern that male migration may violate the exclusion restriction. In particular, if economic shocks promote differential migration among men depending on village polygamy status, this might affect conflict. For example, good economic shocks could differentially reduce out-migration of young men, increasing both the supply of men to the labor market and the demand for brides. The latter would could confound my supply-side mechanism, while the former represents a non-marriage channel by which the instrument could plausibly affect the outcome, violating the exclusion restriction. While it is not immediately obvious why any such effect would be differential by polygamy status, it is worth investigating nonetheless.

If this is the case, I should observe that the IV negatively affects male migration rates. I test this hypothesis using data from the Nigeria General Household Survey (GHS) panel for the years 2010, 2013, and 2015. I use information on whether individuals originally listed as household members have ever migrated away from the cluster in which they are sampled. As in Section 7.3, I link the survey to the rainfall histories of each cluster. Using this data, I estimate an individual-level regression of a migration dummy on the rainfall shock interacted with cluster-level polygamy, including fixed effects for cohort and grid-cell. All models use the sample of men born after 1998, so that their aged 12-16 shock is well-defined and the sample is comparable to the DHS.[45]

The results are given in Table D5. Columns (1)-(3) estimate the model using the individual-level shock – the mean rainfall deviation experienced by a man during the ages of 12-16 – while columns (4)-(6) estimate the model using the same cluster-level shock as in Section 7.3. None of the models are statistically significant, suggesting that there is no differential migration response to the shock; this conclusion does not appear to be sensitive to the inclusion of fixed effects or the level at which the shock is defined. There is weak evidence in column (6) that, after accounting for location and cohort fixed effects, out-migration falls slightly in polygamous areas, but these results are not significant.

D.6 Robustness to measures of marriage inequality

In Table D6, I also consider the robustness of the main results to different measurements of marriage market pressure. Throughout the paper, I use the Gini coefficient (column 1). However, similar results obtain when I measure inequality using the coefficient of variation (column 2). In columns 3 and 4, I use the unmarried rate for the whole sample, and for under-30 men, respectively. Lastly, in columns 5 and 6 I use the marriage market share of young men, defined as under 30 or under 35, respectively. Both first-stage and 2SLS estimates are significant in all of these specifications. Moreover, the magnitudes are similar across the different measurements.

D.7 Robustness to demand-side shocks

Positive shocks for adolescent girls between the ages of 12-16 may also represent positive shocks for their brothers of the same cohort. Therefore, the adolescent rainfall measure may also capture demand-side shocks to the marriage market, which would drive up both bride prices and marriage levels. In this case, the first-stage would be underestimated, since demand-side adjustments would likely reduce inequality. In addition, these shocks might also improve the future outside options of young men, reducing their propensity to enter Boko Haram and violating the exclusion restriction and biasing the main estimates downward.

This is unlikely to be the case given that there exist large gender gaps in age at marriage without concomitant gaps in the underlying age structure which drives cluster-level variation in the shock.

[45] The latest year of rainfall data is 2014.

That is, shocks that affect men at ages 12-16 are unlikely to have affects on the demand side of the marriage market given that men marry, on average, at age 25. Still, to eliminate all doubt, in Table D7 I augment the reduced form, first-stage, and 2SLS specifications to include various male-side rainfall shocks interacted with polygamy and defined in four-year intervals from 14-18 to 20-24.[46] In column (1) the main results are included for comparison. In general, inclusion of these shocks weakens the significance but not magnitude of the main effects, by adding noise and increasing the standard errors. Consequently, I include Anderson-Rubin p-values to conduct inference in cases in which the first-stage F-statistic falls below the threshold for weak instruments. Under these p-values, the 2SLS estimates remain positive and significant at 5% all but columns (7) and (8), in which they become significant at 10%.

D.8 Measurement error, endogeneity, and instrument validity

The IV estimate of the causal effect of interest φ is 2-3 times as large as the OLS estimate. In this section, I investigate whether the differences between OLS and IV estimates are likely to be driven by the endogeneity of marriage inequality, measurement error bias, or invalidity of the instrument. I formalize this discussion using a Bayesian estimation procedure developed by DiTraglia and Garcia-Jimeno (2018). In a linear IV model with measurement error governed by signal-to-noise ratio κ, a regressor arbitrarily correlated with the 2nd stage error $\rho_{u\xi^*}$, and a potentially invalid IV with correlation $\rho_{u\zeta}$, the authors show that these three "structural" parameters are constrained by each other and the covariance structure of the data.

The authors use this insight to develop a Bayesian estimation procedure which incorporates the information contained in pre-specified researcher beliefs (restrictions on parameter values) for set-identification. In practice, this means that once a researcher imposes beliefs about the likely range of measurement error κ and regressor endogeneity $\rho_{u\xi^*}$ on the reduced-form correlations in the data, one can determine the range of causal effects and instrument correlation contained in the identified set, denoted Θ. Although this cannot definitively prove an instrument valid, it can provide support that instrument validity is consistent with the ex-ante beliefs of the researcher.

Following the discussion of Section 6.4, I impose the beliefs that the model contains moderate classical measurement error and negative selection bias. That is, the endogeneity of marriage inequality in the second-stage is likely to be the result of a negative correlation between G_{jr} and u_{jr}. Negative selection on observables supports this as a reasonable belief. For example, Figure D5 demonstrates

[46] These shocks are defined identically to those used throughout the paper, except that they use the male-side data.

that clusters with greater marriage inequality are richer, more southern, and less muslim–all key risk factors for Boko Haram, which is localized in poor, Northern, muslim communities. If these relationships also extend to unobservables, then negative selection is likely. For a moderate level of measurement error, I use $\kappa > 0.8$.[47] I allow negative selection through the restriction $-0.9 < \rho_{u\xi*} < 0$, which rules out only implausibly large degrees of endogeneity.

The results of the Bayesian estimation procedure are presented in Table D8. Panel I shows estimation summary statistics including the OLS and IV results, as well as L, the lower bound on κ (upper bound on measurement error). Panel II shows the results of inference on the identified set Θ. The first column shows share of posterior draws of parameters resulting in an empty identified set, while the second shows share of draws yielding an identified set that contains $\rho_{u\zeta} = 0$ (a valid instrument). The third and fourth columns show Bayesian 90% credible sets for $\rho_{u\zeta}$ and φ, respectively. Panel III shows posterior medians and 90% credible sets for the partially identified parameters φ and $\rho_{u\zeta}$, estimated under a uniform prior. The bottom row of the table considers the belief that $\kappa > 0.8$, while the second row allows $\kappa < 0.8$.

The results are generally supportive of a valid instrument. Regardless of the measurement error allowed, the probability Θ contains a valid instrument, given the beliefs, ranges from 0.85-0.9. In addition, the credible sets for $\rho_{u\zeta}$ are tightly centered around zero in panels II and III. Finally, the partially-identified estimate of $\rho_{u\zeta}$ is -0.02 or -0.04, very close to zero. Inference on φ depends more on measurement error. Under the $\kappa > 0.8$ belief, the credible sets for φ in Panels II and III exclude zero, ruling out a negative causal effect. In fact, the posterior median for φ is 2.04, suggesting that under these beliefs the IV strategy slightly underestimates the implied causal effect. Unsurprisingly, measurement error beliefs imply more uncertainty around φ. Allowing $\kappa < 0.8$ yields much wider credible sets that include zero. We can conclude that while instrument validity is robust, reliable inference on φ depends on the belief of modest measurement error.

As an additional sensitivity test, I re-run the estimation routine for different beliefs on the extent of endogeneity. Greater negative selection implies a greater gap between the IV and OLS results and therefore is more likely to be consistent with a valid IV; therefore, it is worthwhile to consider how small beliefs over negative selection would have to be for a valid instrument to be excluded from the identified set. To do this, I vary the lower-bound of the beliefs on $\rho_{u\xi^*}$ from -0.9 to -0.05, while still maintaining negative selection and $\kappa > 0.8$.

The results are presented in Figure D6. Panel A shows how the φ and 90% credible sets estimated under a uniform prior vary as the lower-bound on $\rho_{u\xi^*}$ rises. The two dashed lines indicate the IV

[47] That is, at most 20% of the variation in the marriage Gini coefficient is due to measurement error.

and OLS estimates. As expected, φ falls as the endogeneity belief is further restricted. However, for all values of the lower-bound, the 90% credible set around the causal effect not only excludes zero but also the OLS value. This suggests that as long as the belief on negative selection is not too small, the IV estimate is reasonable. Panels B and C directly investigate instrument validity. Panel B shows $\mathbb{P}(\text{Valid})$, the probability that the identified set contains $\rho_{u\zeta} = 0$ as the lower bound varies. This probability falls monotonically as expected, but also demonstrates that even at relatively low levels of selection, such as $-0.3 < \rho_{u\zeta^*} < 0$, the probability of a valid instrument is still above 0.6. Panel C shows the estimate of $\rho_{u\zeta}$ with 90% credible sets under a uniform prior. The dashed line indicates the valid instrument belief. For all values of the lower bound, the estimate is always near zero, and the 90% credible set is centered around or near zero. This suggests that for any belief about negative selection, the posterior median is very near a valid instrument.

D.9 Size of the marriage market

In the empirical exercise, the use of the DHS village as marriage market unit assumes strict endogamy – that all marriages occur within a village. However, cross-village marriages are certainly possible in this setting; Mberu (2005) finds that marriage migration is common in Nigeria overall, though general migration rates are low among muslims and in Northern Nigeria in particular. This effect would likely lead me to underestimate the main reduced form and first stage results. Under autarky, men cannot search for brides outside the village, and thus absorb the full impact of the negative supply shock induced by good adolescent rainfall, whereas trade acts as a release valve to absorb supply shocks. If I wrongly assume autarky, then, at the village level, men who can in fact search in neighboring villages should do so at the margin, partially offsetting the effects on bride price, marriage inequality, and conflict. The implication is that if we aggregate the data to larger units, men are less likely to be able to search outside, so we should expect to see larger effects for these aggregate units.

To test this hypothesis, I cluster DHS villages into larger groups. In particular, I employ a k-means clustering algorithm to identify geographically proximate clusters of villages based on latitude and longitude, [48] setting the number of clusters K to be such that there are, on average, 3 villages per cluster.[49] I then re-estimate the main OLS, reduced form, and first stage regressions. The results are given in Table D9. Columns (1) and (2) do so for the attacks outcome variable, while columns (3) and

[48] Intuitively, this algorithm identifies the partition of the set of villages into K groups that minimizes the sum of squared distances within each group k in the partition.

[49] This choice must balance the loss in power against the desire to test the aggregation effect.

(4) use Boko Haram-related deaths. Overall, the results are not substantially changed from the main results in Table 2, columns (2) and (3). The first stage effect sizes are nearly identical (12.8 vs. 10.3 with all controls included). The OLS and reduced form coefficients increase in magnitude, but this difference is driven primarily by the higher baseline level of conflict within a larger aggregate unit. We can conclude that between-village marriages are not common enough to meaningfully affect the results, and that the assumption of autarky is reasonably innocuous.

D.10 Recruitment and attacks

A more controversial assumption is that Boko Haram activity – measured by attacks and fatalities throughout this paper – is a reasonable proxy for recruitment, which is the object of interest in the theoretical model. Lack of data on recruitment makes directly testing this assumption impossible. However, indirect tests indicate that it is not implausible.

The most extensive study on islamist recruitment in Africa comes from the (UNDP, 2017). The UNDP survey roughly 130 ex-Boko Haram recruits in Nigeria, and provide some information on their demographic characteristics. Overall, these ex-combatants are overwhelmingly natives of the four states that form the core of Boko Haram's territory – Borno, Adamawa, Yobe, and Gombe. However, Figure A2 indicates that while Boko Haram concentrates violent activity in Nigeria's northeast, the group is more than capable of attacking remote targets outside its core territories, consistent with reports of coordinated attacks as far afield as Kano, Sokoto, and Abuja (Weeraratne 2017). The ability of Boko Haram to coordinate attacks far from the areas where they draw their fighters suggests a degree of mobility that calls into question the assumption of violence as a proxy for recruitment.

However, an additional observation motivates an indirect test. As observers frequently remark, the types of attacks Boko Haram is capable of launching vary by location – suicide attacks, IEDs, and other bombings are more likely to strike large targets, urban areas, and areas far from Boko Haram's core territories. Mobility is less likely to be a concern for small arms-based attacks, such as raids on villages or engagements with security forces, or mass abductions. Using event-level data, I show in Table D10 that the spatial pattern of attacks indeed differs by attack type. In particular, violence against civilians (column 1) is more common in each of the core states, are are kidnappings. In contrast, suicide attacks, bombings, and "remote violence" (e.g., IEDs) are all more common outside of the group's core areas. The results suggest that, since these various types of bombings exhibit greater mobility, it is exactly these cases in which attacks are likely to be a poor proxy for recruitment.

Therefore, I consider an additional robustness test in which I drop all events that fall into each of these categories and re-run the main estimation. This is undoubtedly a conservative test, since not all bombings occur far afield and we will therefore drop events that do proxy well for recruitment. The results are given in Table D11, with the main effect for all attacks re-printed in column (1) for reference. Using only violence against civilians (column 2) – which is highly spatially correlated with areas of recruitment– is still significant though unsurprisingly smaller. Dropping remote violence (column 3), suicide attacks (column 4), and all bombings (column 5) do not materially change the magnitude or significance of the reduced form or 2SLS results. Thus, the key result obtains among the set of attacks that are likely to be a good proxy for recruitment.

D.11 Differential population structure

The shock s_{jr} is an average of the economic (rainfall) conditions experienced by the women of village j in their adolescent period. As such, it is a cohort-weighted average of rainfall histories for each female cohort c in village j, where the weights ω_{cj} are equal to the cohort c's share of the female population of village j. Therefore, the identifying variation across villages comes from variation in rainfall histories, as well as variation in the age profile of the female population. This is not inherently problematic. However, issues arise if there are systematic differences in cohort shares ω_{cj} by polygamy status. In particular, if polygamous villages have faster population growth, their shocks may be disproportionately weighted toward younger cohorts. Therefore, the coefficient on the interaction term might be driven not by the underlying hypothesized mechanism, but by the fact that polygamous villages have more recent shocks, which affect outcomes directly and outside of the model.

This is only a concern if the female age structure is systematically different in polygamous and monogamous villages. I check this informally by plotting, for each year of the DHS, the share of each cohort in the total sample by cohort year and polygamy status. Figure D7 presents the results. While the polygamy shares do appear noisier, the two series generally follow the same trends. There is no evidence to suggest that more recent cohorts have systematically larger shares because of faster population growth in polygamous villages.

To formally dispel any lingering concerns, I consider a test whereby I replace the village-level shocks s_j with a measure $\bar{s}_j$ that holds cohort-shares constant at the state level, that is

$$\bar{s}_j = \sum_{c=1}^{C} s_{cj}\bar{\omega}_{cs}$$

Where $\bar{\omega}_{cs}$ is the total population share of cohort c in state s. In this way, potential differences in the age structure of the female between polygamous and non-polygamous villages are held constant within Nigerian states. However, each village still retains its rainfall history. Furthermore, conditional on state fixed effects, variation in $\bar{s}_j$ between DHS villages arises exclusively from idiosyncratic differences in rainfall history rather than differences in age structure. This should allay any concerns that differential age structure is driving the results.

Table D12 presents the results. The first stage results remain significant at the 1% level, with F-statistics on the excluded instrument above 10 in all specifications and coefficients very similar in magnitude to those in Table 3. Instrumental variables estimates are also similar to the original results; they are significant at 5% in the two specifications without state fixed effects, and significant at 10% in the specifications with state fixed effects. Reduced form coefficients are significant at 1 or 5% for all but the final specification in column (4).

D.12 Inference

Young (2019) demonstrates that instrumental variables estimates suffer from inferential issues in finite samples. In the presence of non-iid error processes, standard statistical significance tests may be of incorrect size, leading to over-rejection of nulls, even when the cluster structure of errors is accounted for. Resampling-based inference methods – particularly the bootstrap and jackknife – bring rejection probabilities closer to nominal size and perform better than clustered errors in Monte Carlo simulations. I re-estimate the main regressions with traditional bootstrap, wild bootstrap (Wu 1986, Davidson and MacKinnon 2010), and jackknife.

The results are given in Table D13, which presents p-values and confidence intervals in brackets for each method of inference. Perhaps consistent with Young's findings, these methods tend to reduce the significance of the results, with the exception of the wild bootstrap, which actually increases the significance relative to clustered errors. Still, in the reproduction of the main specification of Table 2, the results remain significant at 5% for each method. Further, the results remain significant at 5 or 10% for most specifications. Only the specifications with state fixed effects are no longer significant under jackknife and traditional bootstrap.

Figure D1: Main results, moved by age 16 sample

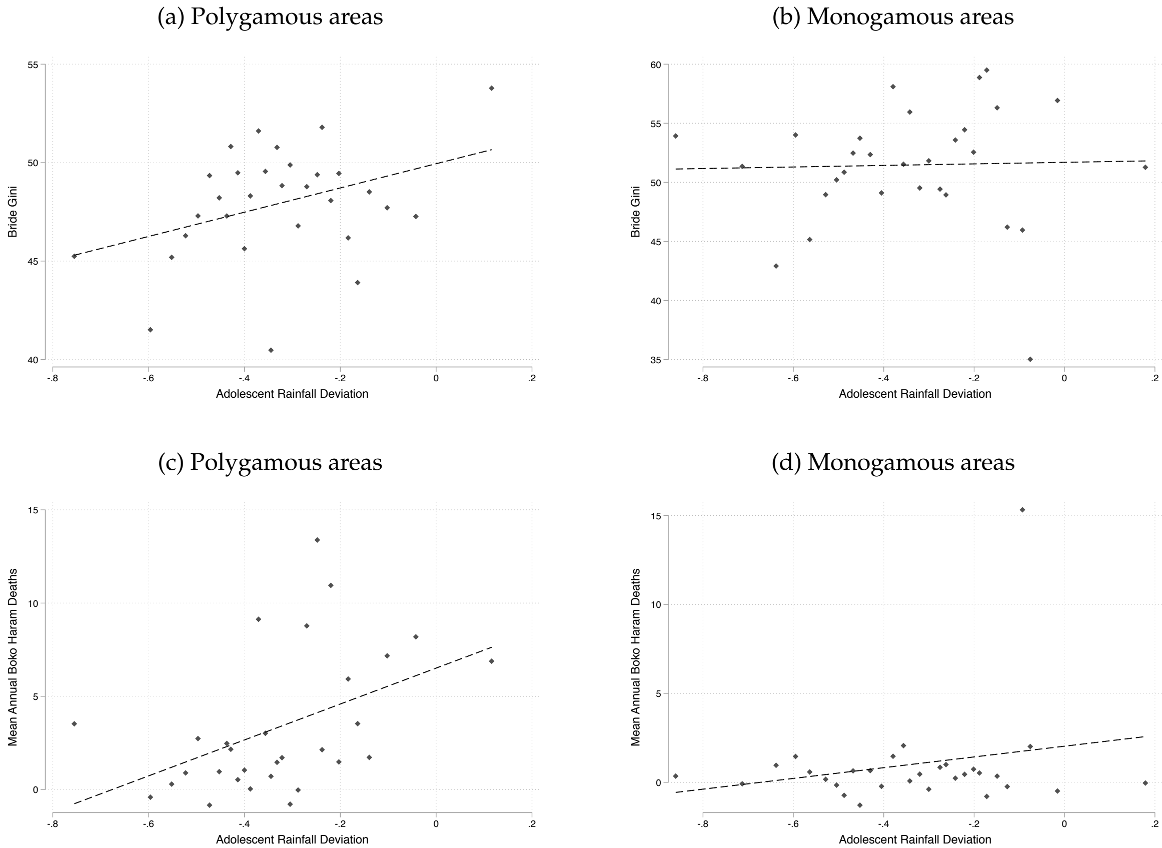

Note: Figure shows the residualized relationship (partial correlation) between the cluster-level mean annual Boko Haram-related fatalities and the mean annual rainfall for women between the ages of 12-16. Scatterplots are binned using 30 quantiles of the distribution of the bride Gini. Panel A estimates the relationship for the subsample of polygamous clusters while Panel B does the same for monogamous clusters. Shocks are defined on sample of women who lived in the cluster in which they are surveyed as of age 16. All plots include DHS round and ethnic homeland fixed effects.

Figure D2: Main results, moved by age 12 sample

(a) Polygamous areas

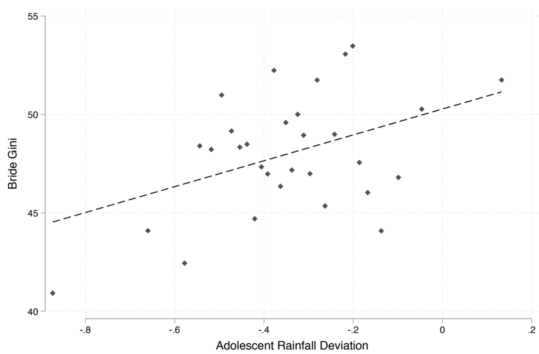

(b) Monogamous areas

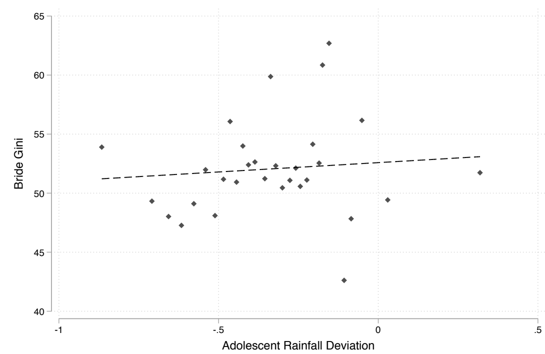

(c) Polygamous areas

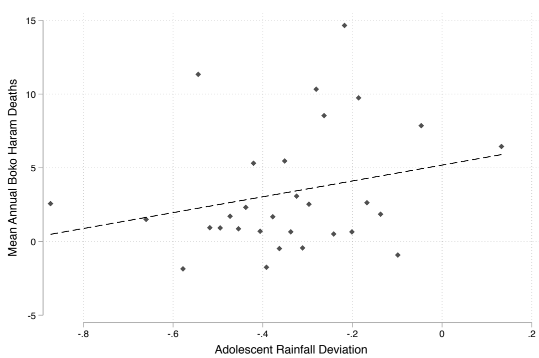

(d) Monogamous areas

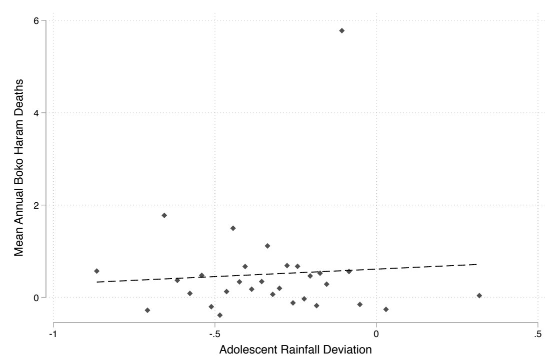

Note: Figure shows the residualized relationship (partial correlation) between the cluster-level mean annual Boko Haram-related fatalities and the mean annual rainfall for women between the ages of 12-16. Scatterplots are binned using 30 quantiles of the distribution of the bride Gini. Panel A estimates the relationship for the subsample of polygamous clusters while Panel B does the same for monogamous clusters. Shocks are defined on sample of women who lived in the cluster in which they are surveyed as of age 12. All plots include DHS round and ethnic homeland fixed effects.

Figure D3: Adolescent rainfall and marriage hazard by age and polygamy status, moved by age 16 sample

(a) Polygamous areas

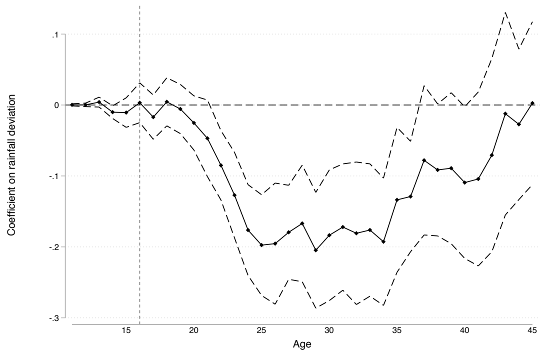

(b) Monogamous areas

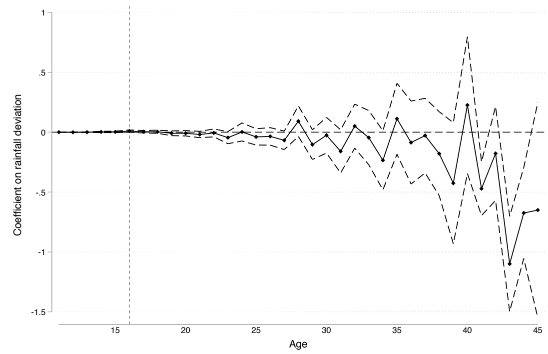

Note: Figure shows coefficients from an individual-level hazard regression of a marriage dummy on mean annual rainfall deviation from ages 12-16 interacted with age dummies, controlling for grid-cell and cohort fixed effects in a sample of women aged 10-45. 95% confidence bands, shown in dashed lines, are calculated using standard errors clustered at the grid-cell level. Panel A conducts the estimation on the subsample of polygamous clusters while Panel B does the same for monogamous clusters. Sample is restricted to women who lived in the cluster in which they are surveyed as of age 16. Vertical line indicates age 16.

Figure D4: Adolescent rainfall and marriage hazard by age and polygamy status, moved by age 12 sample

(a) Polygamous areas

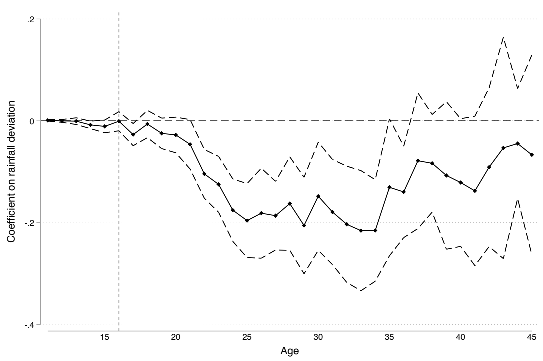

(b) Monogamous areas

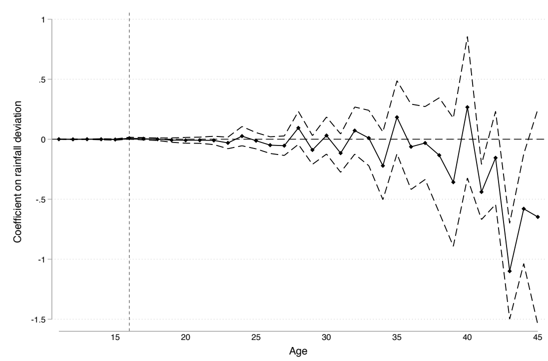

Note: Figure shows coefficients from an individual-level hazard regression of a marriage dummy on mean annual rainfall deviation from ages 12-16 interacted with age dummies, controlling for grid-cell and cohort fixed effects in a sample of women aged 10-45. 95% confidence bands, shown in dashed lines, are calculated using standard errors clustered at the grid-cell level. Panel A conducts the estimation on the subsample of polygamous clusters while Panel B does the same for monogamous clusters. Sample is restricted to women who lived in the cluster in which they are surveyed as of age 12. Vertical line indicates age 16.

Figure D5: Negative selection

(a) Wealth index

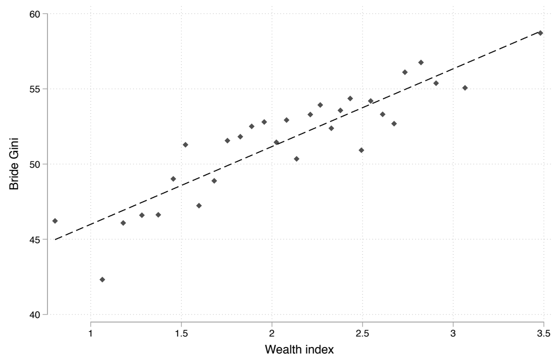

(b) Latitude

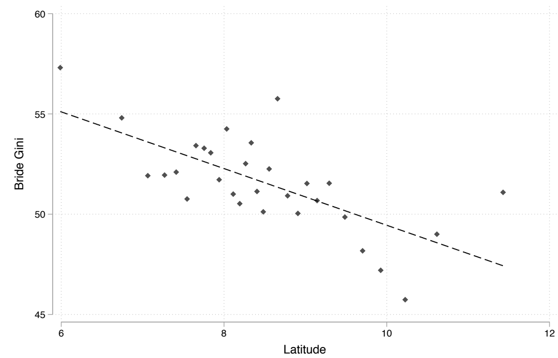

(c) Islam

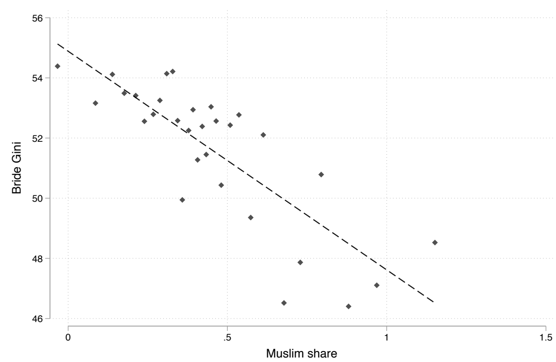

Note: Figure shows the residualized relationship (partial correlation) between the cluster-level mean marriage inequality and average wealth index (Panel A), cluster latitude (Panel B), or muslim share (Panel C). Scatterplots are binned using 30 quantiles of the distribution of the relevant variable. Panels A and B include DHS round and controls for slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Panel C includes all of the above controls except for share muslim.

Figure D6: Sensitivity to negative selection

(a) IV estimate

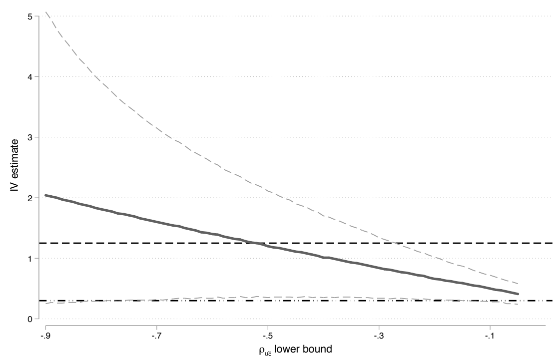

(b) Probability of valid instrument

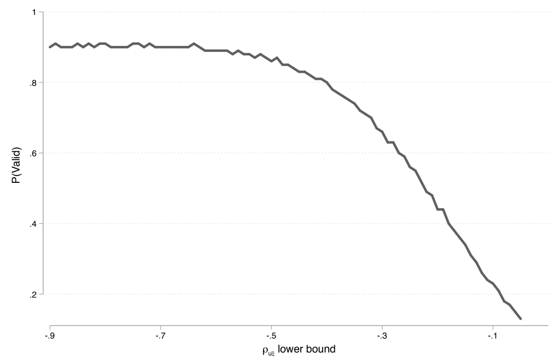

(c) Instrument validity estimate

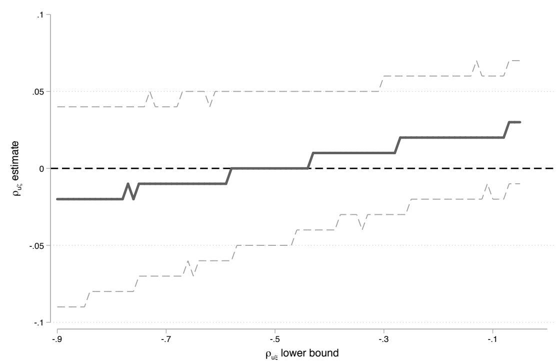

Note: Figure plots results from the Bayesian estimation procedure developed by DiTraglia and Garcia-Jimeno (2018). Dependent variable is average annual number of Boko Haram fatalities within 20 km of the DHS cluster. All specifications include controls for adolescent rainfall shock, polygamy indicator, slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. The plots demonstrate changes in relevant parameter estimates as the lower bound of $\rho_{u\zeta^*}$ varies from -0.9 to -0.05, holding the upper bound fixed at zero and assuming measurement error is classical with $\kappa > 0.8$. Panel A plots the IV estimate and 90% credible intervals taken from column (10) of Table D8 at each value of the lower bound on $\rho_{u\zeta^*}$. Dashed line indicates the frequentist IV estimate, while dash-dotted line indicates OLS. Panel B plots $\mathbb{P}(\text{Valid})$ from column (6) of Table D8 at each value of the lower bound on $\rho_{u\zeta^*}$. Panel C plots $\rho_{u\zeta}$ with 90% credible intervals from column (9) of Table D8 at each value of the lower bound on $\rho_{u\zeta^*}$. See Table D8 for more details.

Figure D7: Cohort composition of female sample by polygamy status and survey year

(a) DHS 2003

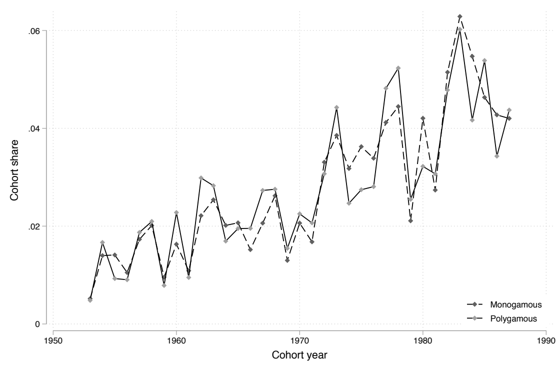

(b) DHS 2008

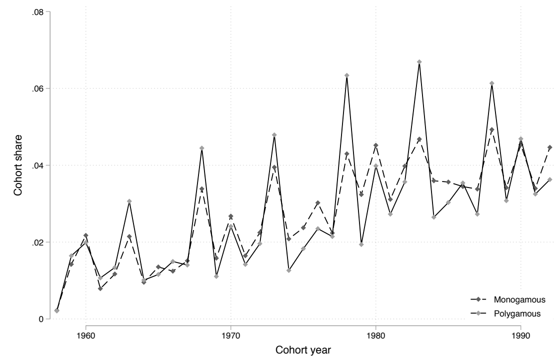

(c) DHS 2013

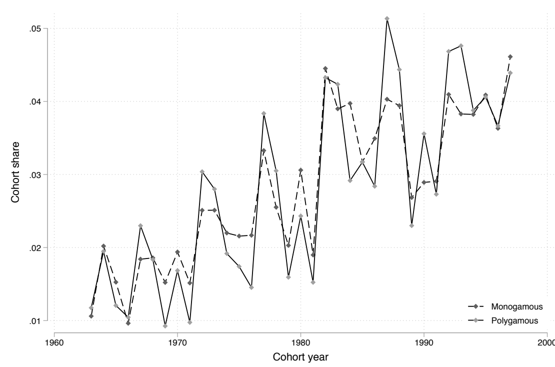

Note: Figure plots the cohort shares for each cohort by polygamy status and DHS round. For each DHS round, I estimate the population shares of each cohort in the monogamous and polygamous subsamples and plot them by cohort.

Table D1: The effect of marriage inequality on violence by geographic area

Dependent variable	Boko Haram deaths					
Sample	North (1)	South (2)	Rural (3)	Urban (4)	Poor (5)	Rich (6)
Panel A: OLS						
Bride Gini	0.640	0.000	0.041	0.265*	0.060	0.055
	(0.388)	(0.000)	(0.035)	(0.148)	(0.045)	(0.088)
R^2	0.452	0.311	0.356	0.706	0.315	0.828
Panel B: Reduced form						
Adolescent rainfall deviation × Polygamous	13.483	-0.082	5.382*	-12.366	10.255*	-3.094
	(29.601)	(0.072)	(3.206)	(12.256)	(5.526)	(6.818)
R^2	0.443	0.313	0.361	0.706	0.319	0.828
Dependent variable	Bride Gini					
Panel C: First-stage						
Adolescent rainfall deviation × Polygamous	21.274***	2.333	14.934***	-3.321	13.143**	2.501
	(6.173)	(3.658)	(4.463)	(5.257)	(5.908)	(3.524)
R^2	0.624	0.531	0.624	0.629	0.620	0.597
Observations	954	814	1125	643	884	884
Controls	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. Outcome variable is either the average annual number of Boko Haram related fatalities within 20 km of the DHS cluster (Panels A, B) or the Gini coefficient of wives in the cluster (Panel C). Tribe FE indicate ethnic homeland fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Poor and rich samples defined as clusters below and above the median of the cluster-level mean wealth distribution, respectively. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D2: The effect of marriage inequality on violence: robustness to polygamy threshold

Dependent variable	Boko Haram deaths						
Polygamy threshold	5%	10%	15%	20%	Male rate	Any female	Female rate
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Reduced form							
Adolescent rainfall deviation × Polygamous	20.169**	12.642**	9.537	15.648**	0.374**	11.561	0.414**
	(8.635)	(5.238)	(5.785)	(6.977)	(0.190)	(9.778)	(0.193)
R^2	0.078	0.080	0.080	0.078	0.080	0.076	0.084
Panel B: Two-stage least squares							
Bride Gini	1.621**	1.437**	1.067*	1.701**	1.232*	1.324	1.751*
	(0.723)	(0.608)	(0.626)	(0.835)	(0.645)	(1.292)	(0.921)
Dependent variable	Bride Gini						
Panel C: First-stage							
Adolescent rainfall deviation × Polygamous	12.443***	8.794***	8.942***	9.201***	0.304***	8.730**	0.236***
	(2.681)	(2.441)	(2.681)	(2.637)	(0.081)	(4.207)	(0.067)
R^2	0.523	0.528	0.531	0.530	0.532	0.515	0.519
Kleibergen-Paap F-statistic	21.536	12.982	11.127	12.172	14.194	4.307	12.354
Observations	1768	1768	1768	1768	1768	1768	1768

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. Outcome variable is either the average annual number of Boko Haram related fatalities within 20 km of the DHS cluster (Panels A, B) or the Gini coefficient of wives in the cluster (Panel C). Polygamy threshold indicates the threshold cluster-level male polygamy share used to classify the cluster as polygamous. In column (5), "male rate" indicates that the raw male polygamy rate is used. In column (6), the any-female criteria is used, while in column (7) the female polygamy rate is used. All specifications include controls and DHS round fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D3: The effect of marriage inequality on violence: robustness to conflict data

Data source	ACLED (1)	GTD (2)	UCDP (3)	SCAD (4)
Panel A: OLS				
Bride Gini	0.477*	0.439*	0.367	0.311*
	(0.270)	(0.232)	(0.240)	(0.174)
R^2	0.088	0.095	0.065	0.088
Panel B: Reduced form				
Adolescent rainfall deviation × Polygamous	20.086**	19.159**	14.617**	15.995***
	(8.576)	(7.933)	(6.871)	(5.962)
R^2	0.078	0.086	0.059	0.084
Panel C: Two-stage least squares				
Bride Gini	1.565**	1.493**	1.139**	1.246**
	(0.693)	(0.640)	(0.553)	(0.502)
Dependent variable	Bride Gini			
Panel D: First-stage				
Adolescent rainfall deviation × Polygamous	12.833***			
	(2.689)			
R^2	0.524			
Kleibergen-Paap F-statistic	22.776			
Observations	1768	1768	1768	1768

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. In Panels A, B, and C, outcome variable is the average annual number of Boko Haram fatalities within 20 km of the DHS cluster. In Panel D, outcome variable is the Gini coefficient of wives in the cluster. Different data sources are indicated in table header. ACLED is Armed Conflict Location and Event Data, GTD is the Global Terrorism Database, UCDP is the Uppsala Conflict Data Program, and SCAD is the Social Conflict Analysis Database. All regressions include DHS round fixed effects and a full set of controls. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels.
** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D4: The effect of adolescent rainfall on marriage supply, robustness to sample

Outcome	Age of marriage				Child marriage			
Lived in cluster as of	age 16		age 12		age 16		age 12	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Adolescent rainfall	0.683***	0.662***	0.761***	0.747***	-0.078**	-0.033	-0.112***	-0.080*
	(0.100)	(0.108)	(0.125)	(0.133)	(0.033)	(0.035)	(0.038)	(0.043)
Observations	16205	16205	11506	11506	4719	4719	3663	3663
R^2	0.250	0.265	0.258	0.278	0.350	0.475	0.248	0.409
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell FE	No	Yes	No	Yes	No	No	No	No
State FE	Yes	No	Yes	No	No	Yes	No	Yes

Standard errors are clustered at the grid-cell level. Outcome variable is either age of marriage, or a child marriage indicator, given in table header. Independent variable is average annual rainfall deviation in the years between age 12 and 16. Sample is a repeated cross-section of individual women between ages 15 and 49. Fixed effects are for DHS round, cohort, state, and/or grid-cell, as indicated. Controls are current rainfall deviation, a quadratic in age, and dummies for Muslim and Hausa. Age of marriage includes the entire sample of ever-married women. For child marriage, the sample includes only women below age 18 and the outcome variable is a dummy marriage indicator. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D5: Robustness to male migration

Level of shock	Individual			Cluster		
	(1)	(2)	(3)	(4)	(5)	(6)
Adolescent rainfall × Polygamous	0.044	0.021	0.030	0.015	0.022	-0.089
	(0.030)	(0.029)	(0.028)	(0.061)	(0.051)	(0.063)
Observations	8596	8596	8596	8596	8596	8596
R^2	0.006	0.097	0.154	0.007	0.099	0.154
Cohort FE	No	Yes	Yes	No	Yes	Yes
Grid-cell FE	No	No	Yes	No	No	Yes

Standard errors, in parentheses, are clustered at the grid-cell level. Dependent variable is an indicator if the individual ever migrated out of the village. Sample is all men in the GHS across rounds 2010, 2013, and 2015 born before 1999. Individual shocks refer to the average annual rainfall deviation in the years when the man was aged 12-16. Cluster-level shocks refer to the cluster-level mean annual rainfall deviation in the years when the women in the cluster were aged 12-16, which is the same definition as is used throughout the paper. ** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D6: The effect of marriage inequality on violence: robustness to inequality measure

Endogenous Variable	Bride Gini	Bride CV	Unmarried		Market share	
			Total	under-30	under-30	under-35
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS						
Marriage inequality	0.477*	14.573*	0.534*	0.289	-0.326	-0.113*
	(0.270)	(7.956)	(0.308)	(0.178)	(0.216)	(0.058)
R^2	0.088	0.085	0.094	0.090	0.084	0.077
Panel B: Two-stage least squares						
Marriage inequality	1.565**	47.305**	1.413**	1.271**	-2.177*	-1.361**
	(0.693)	(20.909)	(0.610)	(0.614)	(1.204)	(0.670)
Dependent variable	Marriage inequality					
Panel C: First-stage						
Adolescent rainfall deviation × Polygamous	12.833***	0.425***	14.218***	15.847***	-9.228**	-14.759***
	(2.689)	(0.090)	(2.823)	(4.537)	(3.673)	(4.669)
R^2	0.524	0.482	0.611	0.345	0.206	0.217
Kleibergen-Paap F-statistic	22.776	22.433	25.361	12.200	6.311	9.994
Observations	1768	1768	1768	1766	1768	1768

Standard errors, in parentheses, are clustered at the grid-cell level. In Panels A, and B outcome variable is the average annual number of Boko Haram fatalities within 20 km of the DHS cluster. In Panel C, outcome variable is the measure of marriage inequality indicated in the table header. All regressions include DHS round fixed effects and a full set of controls. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D7: The effect of marriage inequality on violence: robustness to male shocks

Dependent variable	Boko Haram attacks							
Male rainfall shock	None	14-18	15-19	16-20	17-21	18-22	19-23	20-24
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Reduced-form								
Adolescent rainfall deviation × Polygamous	1.713***	1.372*	1.209**	1.355**	1.077*	1.246**	0.877*	1.118*
	(0.553)	(0.723)	(0.612)	(0.584)	(0.547)	(0.596)	(0.526)	(0.607)
R^2	0.376	0.382	0.381	0.382	0.382	0.381	0.382	0.380
Panel B: Two-stage least squares								
Bride Gini	0.145***	0.132*	0.112*	0.130*	0.095*	0.099*	0.085	0.113*
	(0.051)	(0.075)	(0.064)	(0.070)	(0.054)	(0.052)	(0.055)	(0.068)
Anderson-Rubin p-value		0.049	0.040	0.016	0.041	0.030	0.084	0.056
Dependent variable	Marriage inequality							
Panel C: First-stage								
Adolescent rainfall deviation × Polygamous	11.809***	10.361**	10.753***	10.402***	11.371***	12.638***	10.285***	9.897***
	(2.908)	(4.130)	(4.028)	(4.011)	(3.839)	(3.729)	(3.501)	(3.445)
R^2	0.573	0.573	0.573	0.574	0.575	0.575	0.575	0.575
Kleibergen-Paap F-statistic	16.489	6.292	7.125	6.725	8.772	11.488	8.632	8.255
Observations	1768	1768	1768	1768	1768	1768	1768	1768

Standard errors, in parentheses, are clustered at the grid-cell level. In Panels A, and B outcome variable is the average annual number of Boko Haram attacks within 20 km of the DHS cluster. Columns (2)-(8) includes male rainfall shocks interacted with polygamy, where rainfall shocks are defined over the age range indicated in the table header. Column (1) reprints the main specification for reference. All regressions include DHS round fixed effects, ethnic homeland fixed effects, and a full set of controls. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D8: Measurement error, endogeneity, and instrument validity

	(I) Summary Statistics			(II) Inference for Θ				(III) Inference for θ	
	OLS	IV	L	$\mathbb{P}(\varnothing)$	$\mathbb{P}(\text{Valid})$	$\rho_{u\zeta}$	φ	$\rho_{u\zeta}$	φ
Marriage Inequality ($n = 1768$)	0.30	1.25	0.50						
	(0.09)	(0.46)							
$(\kappa, \rho_{u\zeta^*}) \in (0.5, 0.8] \times [-0.9, 0]$				0.00	0.85	$[-0.17, 0.07]$	$[-11.25, 63.52]$	-0.04	2.48
								$[-0.11, 0.04]$	$[0.34, 6.07]$
$(\kappa, \rho_{u\zeta^*}) \in (0.8, 1] \times [-0.9, 0]$				0.00	0.90	$[-0.15, 0.07]$	$[0.00, 9.30]$	-0.02	2.04
								$[-0.09, 0.04]$	$[0.25, 5.07]$

This table reports results from the Bayesian estimation procedure developed by DiTraglia and Garcia-Jimeno (2018). Dependent variable is average annual number of Boko Haram fatalities within 20 km of the DHS cluster. All specifications include controls for adolescent rainfall shock, polygamy indicator, slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. $\rho_{u\zeta}$ is the instrument validity parameter, while β is the causal effect of marriage inequality on Boko Haram. The first column indicates ex-ante researcher beliefs: I assume measurement error is classical and has a signal-to-noise ratio $\kappa > 0.8$, and selection bias is negative, so that $-0.9 < \rho_{u\zeta^*} < 0$. Columns (2) and (3) report the OLS and IV results for reference. Column (4) reports L, the estimated lower bound on κ. $\mathbb{P}(\varnothing)$ is the proportion of posterior draws where the identified set is empty; $\mathbb{P}(\text{Valid})$ is the proportion of draws that contain a valid IV, $\rho_{u\zeta} = 0$. Columns (7)-(8) report 90% credible intervals on the identified set Θ, while Panel III reports posterior medians and 90% credible intervals for partially identified structural parameters θ.

Table D9: The effect of marriage inequality on violence: robustness to marriage market size

Dependent variable	Attacks		Deaths	
	(1)	(2)	(3)	(4)
Panel A: OLS				
Bride Gini	0.049**	0.122**	0.565**	1.517**
	(0.023)	(0.057)	(0.283)	(0.709)
R^2	0.020	0.139	0.029	0.160
Panel B: Reduced form				
Adolescent rainfall deviaton × Polygamous	0.800	2.149**	15.987*	29.228**
	(0.556)	(1.070)	(8.491)	(13.759)
R^2	0.016	0.105	0.029	0.125
Dependent variable	Bride Gini			
Panel C: First-stage				
Adolescent rainfall deviaton × Polygamous	13.035**	10.278***		
	(5.715)	(3.673)		
Controls	No	Yes		
DHS Round FE	Yes	Yes		
R^2	0.074	0.642		
Observations	590	590		

Sample pools all marriage market k-means clusters over two DHS waves (2008 and 2013). Robust standard errors in parentheses. In Panels A, B, and C, outcome variable is given in the table header. In Panel C, outcome variable is the Gini coefficient of wives in the cluster. All regressions include DHS round fixed effects and a full set of controls. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels.
** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D10: Spatial pattern of Boko Haram attack types

Attack type	VAC (1)	Gov't (2)	Abduction (3)	Suicide (4)	Bombing (5)	Remote (6)	Religious (7)
Borno	0.165***	-0.031	0.010	-0.093***	-0.204***	-0.090***	-0.040**
	(0.032)	(0.035)	(0.009)	(0.021)	(0.032)	(0.026)	(0.019)
Yobe	0.108**	0.047	0.001	-0.061**	-0.134***	-0.069**	-0.017
	(0.043)	(0.046)	(0.014)	(0.027)	(0.041)	(0.032)	(0.027)
Adamawa	0.245***	-0.085	0.039*	-0.111***	-0.249***	-0.119***	-0.004
	(0.051)	(0.052)	(0.023)	(0.029)	(0.041)	(0.031)	(0.027)
Gombe	0.170**	0.001	-0.033***	-0.065	-0.129*	-0.091*	-0.032
	(0.085)	(0.089)	(0.009)	(0.052)	(0.073)	(0.047)	(0.040)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-Statistic	8.512***	1.709	11.574***	5.430***	12.148***	3.876***	1.898
R^2	0.068	0.073	0.020	0.058	0.082	0.081	0.053
Observations	2016	2016	2016	2016	2016	2016	2016

Robust standard errors in parentheses. Dependent variable is indicated in the table header. Sample is 2,016 conflict events in Nigeria from 2009-2017 involving Boko Haram. Independent variables are dummy variables for Borno, Yobe, Adamawa, and Gombe states. All models include year fixed effects. F-statistic is from test of joint significance on the state dummies. ** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D11: The effect of marriage inequality on violence: robustness to type of violence

Dependent variable	All attacks (1)	VAC only (2)	- Remote (3)	- Suicide (4)	- Bombings (5)
Panel A: OLS					
Bride Gini	0.042*	0.016	0.036	0.034*	0.028
	(0.025)	(0.010)	(0.022)	(0.020)	(0.017)
R^2	0.078	0.064	0.076	0.076	0.072
Panel B: Reduced form					
Adolescent rainfall deviation × Polygamous	1.800**	0.610**	1.504**	1.535**	1.156**
	(0.804)	(0.306)	(0.687)	(0.682)	(0.551)
R^2	0.067	0.054	0.066	0.066	0.062
Panel C: Two-stage least squares					
Bride Gini	0.140**	0.048*	0.117**	0.120**	0.090**
	(0.064)	(0.025)	(0.055)	(0.055)	(0.044)
Kleibergen-Paap F-statistic	22.776	22.776	22.776	22.776	22.776
Observations	1768	1768	1768	1768	1768

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. Outcome variable is either (1) all attacks, (2) only violence against civilians, (3) attacks minus remote violence, (4) attacks minus suicide attacks, or (5) attacks minus bombings. All regressions include DHS round fixed effects and a full set of controls. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D12: The effect of marriage inequality on violence: robustness to state-level cohort weights

Dependent variable	Boko Haram deaths			
	(1)	(2)	(3)	(4)
Panel A: Reduced form				
Adolescent rainfall (state weights) × Polygamous	25.146**	8.733**	17.306***	7.791*
	(10.674)	(4.215)	(5.352)	(4.014)
R^2	0.079	0.355	0.392	0.437
Panel B: Two-stage least squares				
Bride Gini	1.968**	0.873*	1.407***	0.718*
	(0.876)	(0.464)	(0.499)	(0.392)
Dependent variable	Bride Gini			
Panel C: First-stage				
$\text{shock}_s tate\times$ Polygamous	12.777***	10.006***	12.299***	10.845***
	(2.781)	(3.015)	(3.042)	(3.259)
R^2	0.528	0.553	0.573	0.585
Kleibergen-Paap F-statistic	21.107	11.014	16.348	11.073
Observations	1768	1768	1768	1768
Controls	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes
Tribe FE	No	No	Yes	Yes
State FE	No	Yes	No	Yes

Sample pools all clusters over two DHS waves (2008 and 2013). Standard errors, in parentheses, are clustered at the grid-cell level. Outcome variable is either the average annual number of Boko Haram related fatalities within 20 km of the DHS cluster (Panels A, B) or the Gini coefficient of wives in the cluster (Panel C). Tribe FE indicate ethnic homeland fixed effects. State FE indicate Nigerian state fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D13: The effect of marriage inequality on violence: resampling-based inference

	(1)	(2)	(3)	(4)	(5)
Bride Gini	0.674	1.565	0.994	1.620	0.810
Cluster robust	0.047	0.024	0.040	0.003	0.057
	[0.009,1.339]	[0.207,2.923]	[0.047,1.941]	[0.564,2.676]	[-0.025,1.645]
Wild bootstrap	0.019	0.004	0.007	0.000	0.008
	[0.096,1.310]	[0.343,2.966]	[0.218,2.075]	[0.571,2.850]	[0.148,1.773]
Bootstrap	0.068	0.035	0.165	0.018	0.215
	[-0.050,1.397]	[0.107,3.023]	[-0.408,2.395]	[0.277,2.963]	[-0.471,2.091]
Jackknife	0.056	0.040	0.097	0.018	0.138
	[-0.017,1.364]	[0.071,3.060]	[-0.182,2.169]	[0.284,2.956]	[-0.261,1.882]
Observations	1768	1768	1768	1768	1768
Controls	No	Yes	Yes	Yes	Yes
DHS Round FE	No	Yes	Yes	Yes	Yes
Tribe FE	No	No	No	Yes	Yes
State FE	No	No	Yes	No	Yes

Sample pools all clusters over two DHS waves (2008 and 2013). The first row reprints the 2SLS coefficient estimates from Table 2. In each subsequent row, p-values and 95% confidence intervals are shown for each method of inference. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. Sample size refers to all panels. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E Derivation of civilian male value function

Consider the decision of a young man who is considering entering the civilian or militant market under regime r before observing his type. His expected value of civilian life will be

$$\begin{aligned}
V_r^{civ}(w_r) = \lambda \frac{N^{RY}}{N^Y} &\left[(v^{RY} - w_r)1(v^{RY} \geq w_r) + 1(v^{RY} < w_r)\beta V_r^{search,RY}\right] \\
&+ \lambda \frac{N^{PY}}{N^Y} \left[(v^{PY} - w_r)1(v^{PY} \geq w_r) + 1(v^{PY} < w_r)\beta V_r^{search,PY}\right] \\
&+ (1-\lambda)\beta V_r^{search}
\end{aligned}$$

With probability λ, he meets a match and draws his type. In the first term the man ends up rich, while the second term covers the case in which he is poor. In either case, the equilibrium bride price w_r is above or below his value v^j. If it is above, he obtains marriage value $v^j - w_r$; if below, he continues to search conditional on type. With probability $1-\lambda$, he obtains the continuation value V_r^{search}.

Assume a stationary equilibrium, so that the equilibrium bride price is always the same . Then

$$V_r^{civ}(w_r) = V_r^{search}$$

$$V_r^{search,RY} = (v^{RY} - w_r)1(v^{RY} \geq w_r) + 1(v^{RY} < w_r)\beta V_r^{search,RY}$$

$$V_r^{search,PY} = (v^{PY} - w_r)1(v^{PY} \geq w_r) + 1(v^{PY} < w_r)\beta V_r^{search,PY}$$

Re-writing

$$\begin{aligned}
V_r^{civ}(w_r) &= \lambda \frac{N^{RY}}{N^Y} \left[(v^{RY} - w_r)1(v^{RY} \geq w_r) + 1(v^{RY} < w_r)\beta \frac{(v^{RY} - w_r)1(v^{RY} \geq w_r)}{1 - 1(v^{RY} < w_r)\beta}\right] \\
&+ \lambda \frac{N^{PY}}{N^Y} \left[(v^{PY} - w_r)1(v^{PY} \geq w_r) + 1(v^{PY} < w_r)\beta \frac{(v^{PY} - w_r)1(v^{PY} \geq w_r)}{1 - 1(v^{PY} < w_r)\beta}\right] \\
&+ (1-\lambda)\beta V_r^{civ}(w_r) \\
&= \frac{\lambda}{1-(1-\lambda)\beta} \frac{N^{RY}}{N^Y} \left[(v^{RY} - w_r)1(v^{RY} \geq w_r) + 1(v^{RY} < w_r)\beta \frac{(v^{RY} - w_r)1(v^{RY} \geq w_r)}{1 - 1(v^{RY} < w_r)\beta}\right] \\
&+ \frac{\lambda}{1-(1-\lambda)\beta} \frac{N^{PY}}{N^Y} \left[(v^{PY} - w_r)1(v^{PY} \geq w_r) + 1(v^{PY} < w_r)\beta \frac{(v^{PY} - w_r)1(v^{PY} \geq w_r)}{1 - 1(v^{PY} < w_r)\beta}\right] \\
&= \frac{\lambda}{1-(1-\lambda)\beta} \left[\frac{N^{RY}}{N^Y}(v^{RY} - w_r)1(v^{RY} \geq w_r) + \frac{N^{PY}}{N^Y}(v^{PY} - w_r)1(v^{PY} \geq w_r)\right]
\end{aligned}$$

F Cohort-specific pattern of shocks

An additional falsification test comes in the form of analyzing the cohort-specific pattern of shocks. Recall that the mean adolescent rainfall shock in a village is a cohort-weighted average of shocks. Intuitively, older cohorts should matter less for marriage market outcomes, since adolescent shocks that induced them to delay marriage should diminish in salience as those women eventually enter the marriage market. Below, I show that if the marriage hazard in any one period is not too large, the cohort-specific effect size observed in a cross section at time t will be non-monotonic in the cohort c: shocks for cohorts in the middle of the age distribution should exert the largest impact on the marriage market.

Consider a marriage market with $c \in C = \{0, ..., C\}$ cohorts participating. Time is $t = 1, 2...T$. There are N_c^f women in each cohort that enters, which we will equalize across c and normalize to one for simplicity. Each cohort has a probability of being married in a given period λ_c^f (the marriage hazard); these are assumed constant over time within a cohort (but not necessarily across cohorts). In the data, we observe a cross-section of C cohorts at some time t. In any period t, the probability a female of cohort c remains unmarried will be

$$(1 - \lambda_c^f)^{t-c}$$

With probability λ_c^f she will marry a husband this period. Thus, the total number of cohort c women supplied to the marriage market at t (the flow) will be

$$\lambda_c^f (1 - \lambda_c^f)^{t-c}$$

What we observe in the data is the overall fraction of married (the stock) of each cohort as of time t, not just those entering marriage in this period (the flow). The stock of married women at the end of period t (going into $t+1$) will be the sum of the newly married and previously married (assuming no divorce).

$$S_{c,t} = \lambda_c^f (1 - \lambda_c^f)^{t-c} + (1 - (1 - \lambda_c^f)^{t-c})$$

Lets assume that the positive income shock for cohort c permanently reduces the hazard into marriage λ_c^f. What is the effect of a change in marriage probability on the stock of married women?

$$g(c) = \frac{\partial S_{c,t}}{\partial \lambda_c^f} = (1 - \lambda_c^f)^{t-c} (1 + t - c)$$

First note that this derivative is always positive for any cohort where $t+1 \geq c$, which is true by definition since by time t there can be at most t cohorts, including cohort 0. It is obvious that increasing (reducing) the per-period marriage probability increases (reduces) overall married population for any cohort. What's less obvious is the behavior of this object with respect to c. Which cohort's shock will have the greatest affect on its marriage rate, and by assumption of equal cohort sizes, the marriage market in aggregate.

In the first case, we can rule out any cohort or time effects on marriage probabilities, so $\lambda_c^f = \lambda^f$ for all c. Ignoring for a moment the discrete nature of c and treating it as continuous for tractability, we can show that g should generally be nonmonotonic in c. This non-monotonicity is driven by opposing stock and flow effects. The reduction in marriage probability affects the flow into marriage most for younger cohorts. If we look at $\lambda_c^f(1-\lambda_c^f)^{t-c}$, it is clear that a lower marriage probability both reduces the immediate likelihood of marrying (first term), but also increases the residual pool of unmarried women (second term). For young cohorts the second term is at or near 1, so that the immediate effect is large. For older cohorts, the second term becomes small so that flow effects are near zero. They are intermediate at middling cohorts. In contrast, the reduction in marriage probability affects the youngest cohort relatively less. In the term $(1-(1-\lambda_c^f)^{t-c})$ of $S_{c,t}$ clearly a reduction in λ_c^f always reduces the stock, but for the most recent cohorts this term is smaller. The result is a nonmonotonic relationship.

The marriage rates under which this nonmonotonicity occurs can be bounded. If the maximizer of g with respect to c is within $[0,t]$, then at t we will observe some (potentially interior) cohort that maximizes the effect. The first derivative with respect to c is

$$g'(c) = (1-\lambda)^{t-c}[(c-t-1)\log(1-\lambda)-1]$$

Which is zero whenever

$$(c-t-1)\log(1-\lambda)-1=0$$

$$c^* = \frac{1}{\log(1-\lambda)} + t + 1$$

Note $c^* \leq t$ whenever $\lambda \leq 1-e^{-1} \approx 0.632$. That is, whenever the rate of marriage is sufficiently low. But it also can't be too low, to ensure that the maximizer is in the interval and therefore the function to be nonmonotonic over the range we care about, we also need $c^* \geq 0$, or $\lambda \geq 1-e^{\frac{-1}{t+1}}$. Note that if we observe many cohorts, this requirement will not bind since we are likely to observe the maximizing cohort (as long as $\lambda < 0.632$). The second-order conditions can be used to verify that c^* is a maximum, and so the effect is maximized somewhere within the interval (although the precise location depends on the values of the parameters). The result is a U-shaped relationship between the marginal effect of marriage shocks on marriage levels and the cohort receiving the shock. Put another way, the shocks that most affect the current marriage market quantities are shocks that affected cohorts somewhere in the middle of the age distribution.

I test this prediction in Table F1 by re-estimating the first-stage and reduced form regressions where the cohort-weighted average of shocks is replaced with shocks calculated using only data for each cohort. Cohorts are grouped from ages 15-49 in intervals of five years. All models include the standard controls, as well as polygamy interacted with a quadratic in average female age to control for differential age effects. The results confirm the prediction of non-monotonicity. In the first-stage regressions, only the individual cohort coefficients for groups aged 20-35 (columns (3)-(5)), those in the middle of the age distribution, are significant at the five percent level. The results correspond in the reduced form, though with a slight lag–the coefficients from ages 25-39 (columns (4)-(6)) are large and significant while the remainder are near zero and insignificant. The results also help to further rule out contemporaneous effects of weather on conflict, as the relevant shocks driving the results are those that occurred many years ago.

Table F1: Cohort-specific shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Boko Haram deaths								
Adolescent rainfall deviation × Polygamous	20.345***							
	(6.251)							
Cohort 15-19 × Polygamous		4.377						
		(4.638)						
Cohort 20-24 × Polygamous			1.106					
			(5.353)					
Cohort 25-29 × Polygamous				7.371***				
				(2.319)				
Cohort 30-34 × Polygamous					16.622**			
					(6.482)			
Cohort 35-39 × Polygamous						7.965		
						(4.961)		
Cohort 40-44 × Polygamous							2.762	
							(5.723)	
Cohort 45-49 × Polygamous								-12.729
								(9.204)
Observations	1768	1756	1756	1759	1754	1734	1700	1649
R^2	0.386	0.382	0.382	0.395	0.389	0.382	0.389	0.360
Panel B: Bride Gini								
Adolescent rainfall deviation × Polygamous	12.059***							
	(2.883)							
Cohort 15-19 × Polygamous		3.493*						
		(1.877)						
Cohort 20-24 × Polygamous			4.292**					
			(1.996)					
Cohort 25-29 × Polygamous				5.902***				
				(1.784)				
Cohort 30-34 × Polygamous					4.121***			
					(1.540)			
Cohort 35-39 × Polygamous						2.038		
						(1.470)		
Cohort 40-44 × Polygamous							2.201	
							(1.630)	
Cohort 45-49 × Polygamous								1.442
								(1.825)
Observations	1768	1756	1756	1759	1754	1734	1700	1649
R^2	0.577	0.572	0.571	0.577	0.572	0.575	0.568	0.580
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
DHS Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tribe FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors are clustered at the grid level. Outcome variable is either the cluster-level Bride Gini or the mean annual Boko Haram-related fatalities, as indicated. Adolescent rainfall deviation is the cluster-level mean of individual annual rainfall deviations from the long-run average between the ages of 12-16. Column (1) calculates the adolescent rainfall shock using the full sample of women, while columns (2)-(8) calculate cohort-specific adolescent rainfall shocks. Tribe FE indicate ethnic homeland fixed effects. Controls are slope, current and lagged rainfall deviation, population density in 2005, average monthly temperature in the survey year, the wealth index Gini, average age of men and women, distance to national borders, share Hausa, and share Muslim. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.