

Not Gone With the Wind: Long-Run Impact of Herbicidal Warfare in Vietnam*

Gaku Ito[†] Duc Tran[‡] Yuichiro Yoshida[§]

Revised: January 26, 2026

Abstract

We investigate the legacies of herbicidal warfare on population size and growth in Vietnam between 2001 and 2020. The identification strategy exploits the flight-level military records and the plausibly exogenous hamlet-level variation around the spray-on, direction-change, and spray-off points of fixed-wing aircraft missions, combined with a fuzzy regression kink design. The elasticity estimate suggests that a 10% increase in predicted herbicide exposure during the Vietnam War is persistently associated with a decrease in population *size* of 1.88–2.51% in 2001 and 2.61–3.27% in 2020. We also find that a 10% increase in herbicide exposure is associated with a 0.75 percentage-point decrease in population *growth rates* over the 2001–2020 period. These findings reveal previously unseen, lasting legacies of herbicidal warfare on contemporary health outcomes. The empirical patterns run counter to the conventional wisdom of a rapid recovery or a catch-up growth in post-war societies from wartime destruction, while also contributing to the emerging debate over the socioeconomic and public health consequences of broader chemical exposure.

Word Count: 9,517 words

Keywords: Agent Orange; armed conflict; historical legacies; population; urbanization; Vietnam War

*We are grateful for the insightful comments from Erik Lin-Greenberg, Masaaki Higashijima, Keisuke Kawata, Yoshiaki Kobayashi, Vally Koubi, Adeline Lo, Wakako Maekawa, Junichi Yamazaki, session participants of the 2022 Annual Meeting of the American Political Science Association, the 2022 Annual Meeting of the Japan Society for International Relations, the 2023 Pacific International Politics Conference, and seminar participants at Hiroshima University, Kyoto University, University of Toyama, and the University of Tokyo. Financial support from the Japan Society for the Promotion of Science is gratefully acknowledged (Grant Numbers 23H00039, 24KK0025, and 24K16354). All remaining errors are ours.

[†]Associate Professor, Graduate School of Economics, Osaka Metropolitan University. 3–3–138 Sugimoto, Sumiyoshi, Osaka 558–8585, Japan. Email: gaku@omu.ac.jp.  0000-0002-5362-6918. URL: <https://gaku-ito.github.io>. Corresponding author.

[‡]Assistant Professor, Global Campus Institute and Graduate School of Integrated Sciences for Life, Hiroshima University. 1–4–4 Kagamiyama, Higashi-Hiroshima, Hiroshima 739–8528, Japan. Email: tranduc@hiroshima-u.ac.jp.  0000-0002-3443-6510.

[§]Professor, School of Economics, Kwansei Gakuin University. 1–155 Uegahara Ichiban-cho, Nishinomiya 662–8501, Japan. Email: yuichiro.yoshida@kwansei.ac.jp.  0000-0002-3061-4594.

1 Introduction

Does a temporary shock of large-scale violence leave lasting effects on socioeconomic outcomes? In some instances, exposure to violence persistently depletes human capital and undermines interpersonal and institutional trust (e.g., Grasse, 2023; Licher et al., 2020; Nunn and Wantchekon, 2011). Elsewhere, forms of violence foster political engagement, social cohesion, and altruistic behavior (e.g., Bauer et al., 2016; Berman et al., 2023; Blattman, 2009). By contrast, several studies also highlight a rapid recovery or a catch-up growth of economic and population outcomes in post-war societies, with the impact of wartime destruction remaining short-lived (e.g., Brakman et al., 2004; Davis and Weinstein, 2002, 2008; Miguel and Roland, 2011; but see Bosker et al., 2007).

This article joins the growing literature with new evidence from the herbicide spray in the Vietnam War. While earlier works highlight a powerful recovery from wartime destruction (Miguel and Roland, 2011), recent literature finds persistent effects of U.S. bombing on development and health outcomes in Cambodia (e.g., Lin, 2022), Laos (e.g., Riaño and Caicedo, 2024), and Vietnam (e.g., Palmer et al., 2019; Singhal, 2019). Most closely related to this article, Appau et al. (2021), Le et al. (2022), Yamashita and Trinh (2022), and Vuong (2024) examine the legacies of herbicidal warfare. For example, Appau et al. (2021) underline a persistent negative association between the district-level herbicide exposure and the agricultural productivity, trust levels, and economic production. In a similar vein, Le et al. (2022) document a commune-level positive association between herbicide exposure and the immediate and detrimental prevalence of health disease and mobility disability in the cohort born before the spray mission ended in 1971.

While the existing insights are valuable, a key challenge for causal identification stems from the commonly used, district- or commune-level geographic aggregation. This is unfortunate, as the distribution of herbicide dispersal is partly a function of plausibly exogenous, micro-level factors including climate and wind conditions, local geography, and instantaneous and *uncontrollable* decision of aircraft pilots besides initial mission plans and intended targets

within geographically-small areas. To illustrate, given the typical airspeed of 240–278 km/h for spray flights (Darrow et al., 1969, 26), even a “one-minute” misalignment in the timing of spray-on/off or changes in direction, which can be triggered by unexpected weather or wind conditions as well as ground fire, could shift the dispersal patterns by more than four kilometers in an unexpected way. The district-level or commune-level aggregation, with an average size of 465.31 km² (districts) or 29.6 km² (communes), inevitably masks such micro-level variation that would be valuable for causal identification.¹ An alternative, straightforward identification strategy is to leverage the plausibly exogenous micro-level variation, which is lost in geographically intensive aggregation.

To do so, we combine the fine-grained archival data previously unused at the original, *flight*-level scale with a fuzzy regression kink (RK) design to address the identification challenge arising from the nonrandom nature of historical exposure to herbicide, most notably Agent Orange, during the Vietnam War. Indeed, historical records suggest that Operation Ranch Hand (1962–1971) missions intended to prevent directly damaging densely populated areas and croplands not under Viet Cong control (MACV, 1969), which would invite bias into naive comparisons. To surmount the identification challenge, our RK strategy leverages the *uncontrollable*, natural experimental variation in the hamlet-level exposure to herbicide around the spray-on, direction-change, and spray-off points of fixed-wing aircraft missions. To capture the micro-level variation in herbicide missions, we rely on the flight-level records documented in the Stellman-National Academy of Sciences version of the Herbicide Report System file (S-NAS HERBS, Stellman et al., 2003a,b).

The RK analysis reveals lasting legacies of herbicidal warfare such that even after decades, herbicide exposure during the 1961–1971 period is negatively and systematically associated with population size and population growth in the present day. The elasticity estimate suggests that a 10% increase in herbicide exposure, which is observed within a 1 km distance from the spray start, turn, and end points, is associated with a 2.61–3.27% decrease in

¹ Assuming a circle of an equivalent area, the oft-employed geographic aggregation masks variations within a diameter of $2 \times \sqrt{465.31/\pi} = 24.34$ km (districts) and $2 \times \sqrt{29.6/\pi} = 6.14$ km (communes).

population size in 2020. The analysis also reveals negative associations between the herbicide shock and population *growth rate* in 2001–2020, and runs counter to the conventional wisdom of a powerful recovery from wartime destruction in Vietnam (e.g., [Miguel and Roland, 2011](#); see also, [Davis and Weinstein, 2002, 2008](#)). A series of placebo tests and robustness checks lends further credibility to the findings. While the discrepancy might arise from different units of analysis and the focus of distinct treatments, the previously unseen patterns run counter to the conventional wisdom of a powerful recovery from wartime destruction. Besides the legacies of war, the findings also speak to the growing body of literature on the economic and health effects of pesticide or chemical usage (e.g., [Camacho and Mejía, 2017](#); [Dias et al., 2023](#); [Frank, 2024](#)).² Collectively, this article offers new insights of the lasting legacies of herbicidal warfare beyond the established notion of the “forced-draft urbanization and modernization” in Vietnam ([Huntington, 1968](#), 652). As such, our findings challenge the existing findings on wartime recovery while contributing to the debate over the socioeconomic and public health consequences of broader chemical exposure.

The remainder of this article proceeds as follows. The next section reviews the historical background and related literature. Section 3 details the archival records used for the empirical analysis, followed by the identification strategy and estimation models in Section 4. Section 5 reports the main findings, along with a summary of the sensitivity analysis and robustness checks. Section 6 concludes.

2 Lasting Legacies of Chemical Exposure

The military herbicide operations during the Vietnam War represent the largest episode of chemical-weapon deployment in modern history. Between 1961 and 1971, U.S. forces sprayed more than 20 million gallons of herbicides across South Vietnam. The primary objective was the large-scale defoliation of forests and mangroves to deprive Viet Cong insurgents of natural cover and to improve aerial targeting efficiency. Secondary objectives included the

²We thank an anonymous reviewer for suggesting this body of literature.

destruction of enemy food supplies and the defoliation of military base perimeters, accounting for approximately 9% and 2% of total herbicide dispersal, respectively (IoM, 1995, Stellman and Stellman, 2018). More than 95% of all herbicides were dispersed by the U.S. Air Force under Operation Ranch Hand, primarily through fixed-wing C-123 aircraft (Stellman et al., 2003a). Other military branches employed helicopters, boats, trucks, and hand sprayers, but these accounted for substantially smaller quantities. The composition of herbicides varied by mission type. Agent Orange, contaminated with the highly toxic and persistent dioxin 2,3,7,8-tetrachlorodibenzo-p-dioxin (TCDD), was used most intensively for forest defoliation.

Unlike other herbicide ingredients, TCDD exhibits exceptionally high toxicity. It is chemically stable, resistant to degradation, and therefore persists in the environment for extended periods. Due to its strong lipophilicity, TCDD bioaccumulates in human adipose tissue and biomagnifies within ecosystems. TCDD functions as an endocrine-disrupting chemical through activation of the aryl hydrocarbon receptor, thereby interfering with hormonal regulation and developmental signaling pathways (IARC, 2012). Epidemiological studies of highly exposed populations, most notably the Seveso cohort in Italy, demonstrate associations between elevated TCDD body burdens and impaired reproductive function, adverse pregnancy outcomes, and long-term increases in cancer incidence (Consonni et al., 2008; Mocarelli et al., 2000). Prenatal and early-life exposure has also been linked to fetal growth retardation, congenital abnormalities, and neurodevelopmental impairment, supporting TCDD's classification as a Group 1 human carcinogen (IARC, 2012).

Given its bioaccumulative and persistent nature, a significant body of research has examined the long-term health and socioeconomic consequences of dioxin exposure during the Vietnam War. Early studies focusing primarily on U.S. veterans report associations between herbicide exposure and cardiovascular disease, cancers, diabetes, Parkinson's disease, and adverse reproductive outcomes, although the confidence of the evidence varies considerably (IoM, 1995, 6–7; Stellman et al., 1988). More recent research, situated within the broader literature on the long-run legacies of war, has shifted attention toward civilian populations in

Vietnam. Beyond the short-term wartime displacement emphasized by Huntington (1968), these studies document persistent effects decades after the cessation of hostilities. For example, Yamashita and Trinh (2020) provide evidence linking herbicide exposure to elevated cancer incidence and disability among Vietnamese civilians. Using contemporary census data, Appau et al. (2021) show that district-level herbicide exposure reduced household agricultural productivity in Vietnam. Le et al. (2022) also find that cohorts born before the end of the spray missions in 1971 experienced higher disease prevalence and mobility disability at the commune level.

The literature suggests that military herbicide exposure may generate persistent population decline through three main mechanisms. First, direct biological effects can reduce fertility and increase infant mortality, consistent with evidence from the pesticide literature documenting impaired reproductive outcomes following chemical exposure (Camacho and Mejía, 2017; De Araujo et al., 2023). Even modest reductions in cohort size at birth can cumulate into long-run population losses when exposure is spatially concentrated and persistent. Second, exposure may induce selective migration. Contaminated areas are likely to experience sustained out-migration and reduced in-migration as households respond to health risks, environmental degradation, and weakened economic prospects. Evidence linking herbicide exposure to disability and lower human capital suggests that migration responses may reinforce initial population losses, generating persistent spatial divergence (Yamashita and Trinh, 2022). Third, persistent soil contamination may depress agricultural productivity and farm profitability, reducing the land's economic carrying capacity. Empirical studies documenting declines in crop yields and long-term health among working-age cohorts imply that population decline may arise not only through health and migration channels, but also through structural constraints on local livelihoods (Appau et al., 2021; Le et al., 2022). Reduced productivity weakens income generation and food security, further discouraging settlement and family formation.

Beyond the literature on wartime violence, this study contributes to a broader body of

research on the demographic and health consequences of agrochemical exposure. Although TCDD entered the environment through military herbicide operations rather than agricultural application, its toxicological properties and modes of action closely resemble those of several widely studied pesticides. The comparability between TCDD and agricultural pesticides arises from similarities in their mechanisms of action and toxicological characteristics.³ In this sense, wartime dioxin exposure can be viewed as an extreme and historically unique case of large-scale agrochemical contamination, allowing insights from the pesticide literature to inform the interpretation of its long-run effects. Evidence from non-conflict settings indicates that intensive herbicide and pesticide use has substantial and persistent impacts on population health. Large-scale aerial spraying of glyphosate in Colombia led to sharp increases in medical consultations and miscarriages (Camacho and Mejía, 2017), while rising glyphosate application in Brazil significantly increased infant mortality rates (De Araujo et al., 2023). More broadly, numerous studies document adverse effects of pesticide and fertilizer exposure on infant and child health across diverse contexts (Brainerd and Menon, 2014; Calzada et al., 2023; Frank, 2024; Lai, 2017; Skidmore et al., 2023). These findings indicate that chemical exposure affects not only individual health outcomes but also long-run demographic dynamics. This study situates wartime herbicide deployment within this broader framework, treating it as an extreme case of large-scale chemical exposure.

Understanding the long-run impact of wartime herbicide exposure also requires situating the analysis within Vietnam's post-war historical and institutional context. Following reunification in 1975, Vietnam faced widespread infrastructure destruction, environmental degradation, and severe public health challenges. Although reconstruction began immediately, economic recovery remained slow and uneven until the market-oriented *Đổi Mới* reforms initiated in 1986. While these reforms generated rapid national growth and substantial poverty

³First, both TCDD and several classes of pesticides, particularly organochlorine compounds and certain insecticides and herbicides, act as endocrine-disrupting chemicals, interfering with hormone receptors and transcriptional regulation that govern reproduction, fetal development, metabolism, and neurological function. Second, both exhibit high lipophilicity and resistance to degradation, enabling persistence in soils and sediments, bioaccumulation in human adipose tissue, and biomagnification through food chains. Third, for both, health risks arise primarily through low-dose, chronic exposure rather than acute toxicity.

reduction, recovery trajectories varied markedly across regions. Areas heavily exposed to wartime herbicides often lagged in health outcomes, agricultural recovery, and infrastructure development. Remediation and health monitoring programs targeting dioxin contamination were introduced primarily from the 1990s onward, but their geographic coverage and effectiveness remained limited (Stellman and Stellman, 2018).⁴ Against this backdrop, persistent population deficits in heavily sprayed areas would indicate that the demographic consequences of chemical exposure were not fully offset by post-war growth and policy intervention. This analysis therefore examines whether wartime herbicide deployment left a durable imprint on population dynamics nearly half a century after exposure.

3 Data

Our empirical analysis combines the archival military records of herbicide spray, historical maps, and contemporary geospatial datasets to explore the persistent legacies of herbicidal warfare in Vietnam. To mitigate the endogeneity concerns, we leverage the micro-level variations in spray flights within geographically small areas around the nozzle-on/off and direction-change points of Ranch Hand aircraft, combined with the regression kink design.

3.1 Stellman-National Academy of Sciences HERBS File

To measure the landscape of herbicide spray, we rely on the military archival records of the Stellman-National Academy of Sciences version of the Herbicide Report System file database (S-NAS HERBS, Stellman et al., 2003a,b).⁵ The S-NAS HERBS is a compilation of military

⁴Dung (2023) provides a brief overview of post-war support for Agent Orange victims. Specifically, under Vietnamese government’s “Preferential Treatment of People with Meritorious Services to the Revolution,” veterans who suffer from health issues caused by wartime violence, including exposure to herbicide spraying, and their families are eligible to receive monthly allowances. Civilians affected by Agent Orange are eligible for monthly allowances under the social protection policies for people with disabilities. Agent Orange victims also receive assistance from domestic and international non-government organizations.

⁵We web-scraped the mission records from the Agent Orange Warehouse website. Available at: <http://www.workerveteranhealth.org/milherbs/new/>, accessed August 29, 2021. See Appendix A for details. We discarded 91 errant entries of “fixed-wing aircraft” records with single (not multiple) geocoordinates.

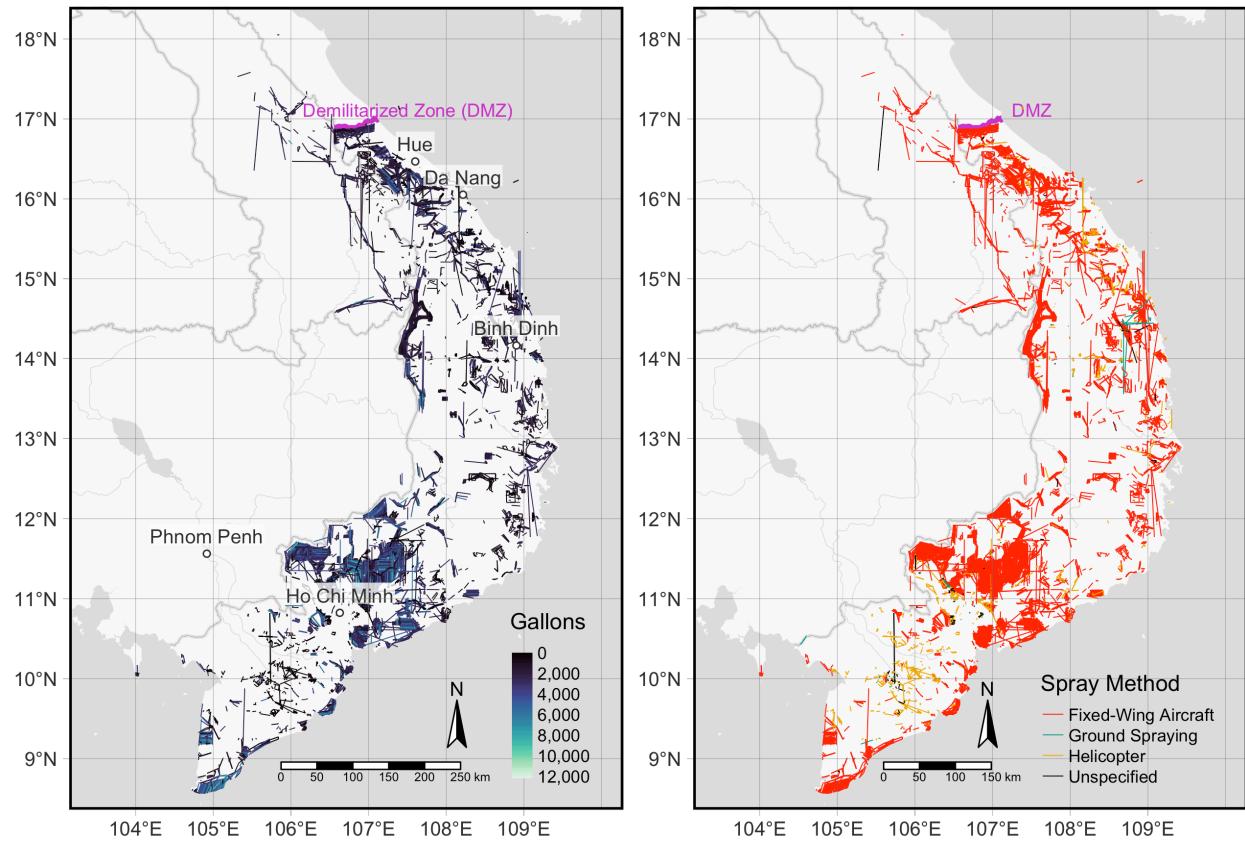
mission records of the HERBS file originally developed by the U.S. Department of Defense with corrections, containing 9,141 reports of the spray missions with several spray methods (i.e., fixed-wing aircraft, helicopters, and ground-spraying) in 1961–1971. A record in the S-NAS-HERBS database corresponds to a single mission with one or more spray paths. Figure 1 shows the geographical and temporal distributions of the mission records. As elaborated in detail below, the following analysis relies on the fixed-wing aircraft missions that dispersed approximately 95% of herbicides (Stellman et al., 2003a, 681–682).

A noteworthy aspect of the database is its geocoding accuracy. In addition to mission dates, spray methods, agent types, and gallonage information, the database documents “the actual flight paths taken by Ranch Hand aircraft as they carried out their spray missions,” including “locations at which the aircraft switched directions or turned off and on their spray nozzles” (Stellman et al., 2003b, 323), providing a detailed landscape of spray missions.⁶

3.2 Treatment and Herbicide Exposure Opportunity Index

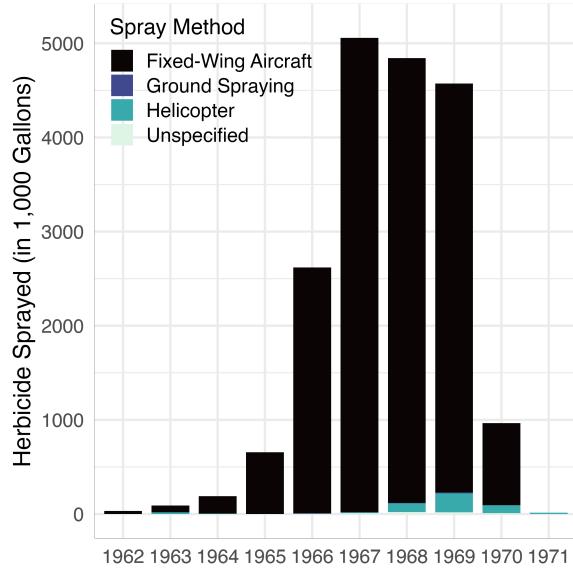
We use hamlets in South Vietnam as the unit of analysis in the following analysis, based on the geocoordinates recorded in the Vietnam Hamlet Evaluation System (HES) Gazetteer Data (Douglass, 2011). The key treatment indicator is the spray-on, direction-change, and spray-off of fixed-wing aircraft spray missions. Thus, we fundamentally estimate the causal effect of herbicide spraying on post-war population growth, measured as the distance gradient from the spray nozzle on-off and direction switching points. Solely for interpretive purposes, we then convert the estimated “distance effect” into a “quantity-based effect.” We do so by constructing a numerical model that maps our distance measure to an expected herbicide exposure, which we refer to as an “exposure opportunity index” (EOI) following the literature (e.g., Stellman et al., 2003b; Stellman and Stellman, 1986). Henceforth, although

⁶To navigate aircraft and record spray-on, direction-change, and spray-off points, the tactical air navigation (TACAN) system distance measuring equipment (DME) was well-developed in Southeast Asia by the 1960s. During mission flights, TACAN/DME continuously offered geographical navigation using the (short) ultra high frequency (UHF) radio range (Rowley, 1975). An essentially same navigation system guided civic aviation until the mid-1990s or before the civic use of the global positioning system (GPS).

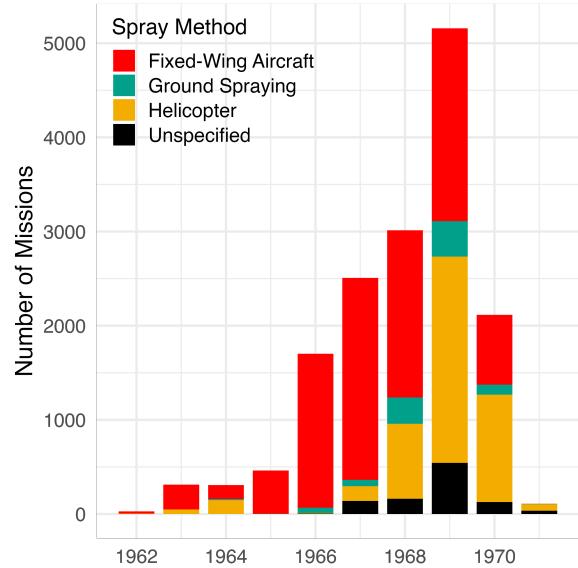


(a) Gallons Sprayed

(b) Spray Methods



(c) Gallons Sprayed by Spray Methods



(d) Number of Missions by Spray Methods

Figure 1: Herbicide Spray Missions

Notes: Segments in maps represent the spray paths. (a) Shading is proportional to the quantity of sprayed herbicide in gallons sprayed by individual spray legs (straight lines without a turn). (b) Colors indicate spray methods. Thick gray lines indicate international borders. Thin gray lines and polygons represent major water bodies. (c), (d) Yearly distribution of herbicide spray (in 1,000 gallons) and spray methods, omitting four helicopter missions in 1961.

our causal estimates rely exclusively on distance and hamlet locations, to facilitate better sense of understanding of the treatment effect, we also report the estimates with hamlet-level opportunities, although not actual levels, of exposure to herbicide.

Specifically, we broadly follow the distance-weighted approach of Stellman and Stellman (1986, 309) and construct an EOI of herbicide spray with the following specification:

$$\text{HERB}_i = \sum_{j \in \mathcal{H}} G_j \cdot e^{-\lambda D_{ij}}, \quad (1)$$

where i indexes hamlets and j spray flight “legs” of the fixed-wing aircraft missions. Following Stellman et al. (2003b, 325), we first split individual flight paths (contiguous lines) into distinct flight legs indicated by the straight lines without turns. As the gallonage information is aggregated at the mission level, we divide and assign the total mission-level herbicide quantities to individual legs in proportion to its spray run length.⁷ G_j indicates the gallons of the herbicides sprayed in leg j , \mathcal{H} is a set of spray flight legs, and D_{ij} represents the geodesic distance between hamlet i and flight leg j . $\lambda = \frac{\ln(2)}{D^{\text{Half}}}$ is an arbitrary parameter, where D^{Half} determines the geodesic distance at which the herbicide exposure decays to the half of a “direct” (zero-distance) hit with $D_{ij} = 0$, with the baseline of $D^{\text{Half}} = 500$ meters.

Two aspects of the exposure opportunity index are worth explanations. First, the index is not a measure of the actual *levels* of exposure to herbicide. Rather, it can be considered a measure of exposure *opportunity*, capturing the likelihood that individual hamlets were exposed based on spray intensity and geographic proximity to flight paths. Admittedly, as with other EOI measures, the absence of direct measurements of biological dose prevents a direct validation of the index against ground truth. Nonetheless, the presented index offers a simple, if inevitably imperfect, proxy for location-based exposure history, and provides an interpretable way to characterize the treatment effect.

Second, the baseline setting of the half-decay distance parameter of $D^{\text{Half}} = 500$ meters

⁷For example, if a mission includes one flight with four legs (three turns) with an equal leg length sprayed 1,000 gallons in total, we assign 250 gallons to each flight leg.

reflects the drift estimates in U.S. military guidance on aerial herbicide applications (Darrow et al., 1969). Specifically, the guidance reports lateral drift for coarse droplets based on droplet fall time multiplied by crosswind speed; and forward (momentum) drift can approximately be estimated in the same manner using aircraft speed instead of crosswinds.⁸ Assuming the same fall time, probable forward drift distance ranges from 2.8–3.2 km for 200 μm droplets to 0.9–1.1 km for 300 μm at the typical aircraft speed of 240–278 km/h. For robustness checks, Appendix C.2 reports the estimates with alternative parameter values of $D^{\text{Half}} = 100$ meters, 250 meters, and 1 km, mimicking smaller and larger drift ranges.

3.3 Outcome: Population Size and Growth Rate

The following analysis focuses on contemporary population size and population growth rate as outcome variables. We rely on the WorldPop data, which provides geographically disaggregated records of population estimates at 100 meter level on an annual basis in the 2001–2020 period.⁹ To construct the population size variable, we first extract the annual records of population counts at the hamlet locations, and take the natural logarithm of the variable after adding 1. The (logged) overall and annual population growth rates, respectively, are measured as $\ln \text{Population}_{2020} - \ln \text{Population}_{2001}$ and $\ln \text{Population}_t - \ln \text{Population}_{t-1}$.

⁸Most of the spray volume of aircraft missions consists of droplets 200 μm or larger (Darrow et al., 1969, 29). 200–300 μm droplets under crosswinds of 3, 6, and 9 mp/h are expected to be laterally displaced by roughly 137, 274, and 411 feet (\approx 42, 84, 125 meters) for 200 μm , and 61, 122, and 183 feet (\approx 19, 37, 56 meters) for 300 μm , respectively.

⁹Available at: <https://www.worldpop.org>, accessed August 6, 2021. Due to errant entries, we exclude the population estimates in 2000 from the analysis. While WorldPop is the sole available dataset to construct the hamlet-level population estimates, we validate the dataset against the national census at the aggregated commune level. The national census is only available for 1999, 2009, and 2019 in Vietnam, with the most disaggregated level remaining the commune level. Comparing WorldPop with the population size in the 2019 census at the commune level, we find a strong positive correlation between the two records (Pearson's $r = 0.853$; Spearman's $\rho = 0.822$), while the discrepancy between the two is only weakly correlated with the census population size ($r = -0.078$; $\rho = -0.070$), suggesting that WorldPop provides a reasonable proxy of the demographic geography both in urban and rural areas.

3.4 Covariate: Key Spray Targets and Related Attributes

To facilitate the empirical analysis, we combine archival sources and several geographical databases to construct three sets of covariates. The first set of covariates measures the proximity to the key spray targets. Operation Ranch Hand involved two primary objectives, forest defoliation and destruction of enemy food supplies, with key targets including “base camps and fire support bases...lines of communication, enemy infiltration routes, and enemy base camps” (IoM, 1995, 85). The mission authorization process also intended to prevent damage to densely populated areas and crops not under Viet Cong (VC) control (MACV, 1969). To measure proximity to the key spray targets, we include the prevalence of VC control in 1967–1969 (HES, McCormick, 2021), geodesic distances to suspected areas of North Vietnam Army (NVA) bases (Enemy Base Area File, BASFA), U.S. Air Force and Navy bases, U.S. Army and Marine troops in 1961–1971 (S-NAS-HERBS), and roads (including trails, *Indochina Atlas*), average hamlet population size in 1967–1969 (HES, Douglass, 2011), and dummy variables for the presence of forest, rice croplands, and slash and burn cultivation (*Indochina Atlas*), with the details described in Appendix A.4.¹⁰

Second, geographic covariates include mean elevation and its standard deviation as a proxy of terrain ruggedness (USGS, 1996), soil suitability for rice cultivation (Zabel et al., 2014), distance to rivers, flow accumulation (Lehner et al., 2008), average precipitation, and average temperature (1970–2000, Fick and Hijmans, 2017). To account for wind conditions potentially influencing drift, we also measure the averages of the u - and v -components (eastward and northward, respectively) of the neutral wind over the mission period (1961–1971, Hersbach et al., 2020).¹¹ Finally, historical covariates are the number of hamlets within 30 kilometers, distance to railways (*Indochina Atlas*) and international borders as of the Viet-

¹⁰Where multiple hamlet-month observations are available, VC control measures the average of a dummy variable indicating VC control. To measure road proximity and cropland presence, we georeferenced and image-processed the maps of *Indochina Atlas* compiled by the Central Intelligence Agency (CIA) in 1970. For the locations of NVA bases, we rely on the Enemy Base Area File (BASFA), July 1, 1967–July 1, 1971.

¹¹The u - and v -components of the neutral wind are measured in meters per second and may take both negative (westward or southward) and positive values (eastward or northward). Together, these components account for both the *direction* and the *magnitude* (speed) of the neutral wind.

nam War period, and the average of the annual minimum distance to aerial bombing drop points in 1965–1971 by USAF (Defense Digital Service, 2016).¹²

4 Identification Strategy: Regression Kink Design

To explore the evidence of causal effects, we rely on a fuzzy regression kink (RK) design that exploits the exogenous fluctuations in herbicide exposure within geographically small areas around the spray start, turn, and end points.¹³ Recall that the S-NAS-HERBS database that documents the “actual flight paths” and the “locations at which the aircraft switched directions or turned off and on their spray nozzles” (Stellman et al., 2003b, 323).¹⁴ The core idea behind our RK strategy is that the *exact* locations at which aircraft makes turns and turned on and off spray nozzles and the herbicide dispersal were, at least partly, *uncontrollable*. Besides initial mission plans and intended targets, the *realized* distribution of herbicide was driven by plausibly exogenous micro-level factors including climate conditions, wind, terrain, and turbulence from the aircraft as well as ground fire hits (IoM, 1995, 86–87).

The RK strategy aligns with the known historical facts. Consistent with the underlying idea, “[t]he responsibility for flying the C-123 during the crucial spraying part of each mission was shared between the pilot and the copilot”: the pilot maintained “control of the switches which started and stopped the spray,” while the copilot was “responsible for handling emergencies” (Buckingham, 1982, 37). As noted in Buckingham (1982, 37), this latter task demanded extreme proficiency, given that “[c]onsidering the low altitude at which Ranch Hand flew, the copilot’s reaction in such an emergency had to be immediate and correct the first time; there would be no chance to rectify a mistake.” As such, instantaneous

¹²The 30 km cutoff reflects the historical facts of Ranch Hand missions and standard error clustering in the regression estimation. See model specification section for details.

¹³As mentioned earlier, our goal is to estimate the causal effect of herbicide spraying on post-war population growth, expressed as the distance gradient from the spray nozzle on-off and direction-switching points. To do this we use the spray nozzle on-off points (and aircraft turning points) as cutoffs in a RK framework, where the running variable is the distance from the spray start, direction-change, and end points.

¹⁴See footnote 6 for the aircraft navigation system in the Vietnam War period. Appendix B investigates potential nonlinearity in the covariate distributions across the kink point, and Appendix C examines model dependence and how model specification influences the main findings.

decisions of aircraft pilots and the disturbing factors jointly determined the actual spray dispersal, which inevitably deviated from the initial mission plans and generated haphazard, natural-experimental variation in the herbicide distribution.¹⁵

Consequently, within geographically small areas around the spray start, turn, and end points, the realized herbicide dispersal generates a discontinuous *slope change*, or a discontinuity in the first derivative, in the distribution of herbicide exposure, which remains uncorrelated with potential confounding forces. Our RK strategy leverages this natural experimental kink in the treatment function, in the absence of kink in the confounder distributions, to derive causal identification.¹⁶ Intuitively, we compare hamlets that were barely covered by spray flights and received direct hits (i.e., “treated” hamlets with greater herbicide exposure due to direct hits) with hamlets that were sufficiently close to but located barely outside of the spray paths and received indirect hits (i.e., “control” hamlets with less exposure due to accidental drifts). Located within geographically small areas, these treated and control hamlets should have similar or smooth (without kink) distributions of geographic, historical, and socioeconomic attributes prior to herbicide exposure across the spray-on/off and direction-change points, which we empirically validate in the following section.

Here, we adopt a RK design rather than a regression discontinuity (RD) design given the probable drifts of herbicide outside the areas with direct hits. We rely on a fuzzy, not

¹⁵To illustrate, approximately 29% of aircraft sorties were intercepted by ground fire in 1966 (IoM, 1995, 86), and the crop damage induced by drift on defoliation missions was greater than the damage by crop destruction missions (National Academy of Sciences, 1974, S-5).

¹⁶In RK designs, the cutoff itself need not be exogenous. Rather, similar to RD designs, identification relies on the smoothness of covariate distributions and the lack of deterministic manipulation of the running variable (smoothness assumption, Card et al., 2015b). Often, policy thresholds are set *non-randomly*; nonetheless, as long as the smoothness assumption holds, the RK designs provide unbiased estimates of the policy impact or the treatment effect. In the current setting, however, the assignment of the cutoff is locally exogenous. This is because the spray-on, spray-off, and direction-switching locations in our data reflect *actual* operational records rather than intended flight plans, thereby strengthening the credibility of the smoothness assumption. Metaphorically, the pilot’s “trembling hand” introduces local randomness into the precise locations of herbicide release, which ensures that there are no systematic discontinuities in potential outcomes at the kink points. Below, we empirically validate the smoothness of the distributions of the key hamlet-level covariates around the cutoff, along with the running variable distribution. Note also that just as the fuzzy RD design, the fuzzy RK design is essentially an IV estimation conducted within the local neighborhood around the cutoff, which allows incomplete manipulation while addressing measurement errors in the treatment and the running variable (Card et al., 2015b, 2467–2469).

sharp, RK design given the unknown slope change parameter of the treatment function at the kink point, as well as probable measurement errors in the treatment and the running variable retrieved from the archival military records (Card et al., 2015*b*, 2464–2467).

4.1 Running Variable

The RK design requires a subset of sample hamlets located within sufficiently small geographic areas around the spray start, turn, and end points. The analysis also requires the corresponding distance measure between hamlet locations and the spray start, turn, and end points as a running variable as well as a treatment-side dummy variable indicating actual spray hits. More specifically, our running variable is the distance from the cutoff, or the spray start, turn, and end points, measured as a continuous real number that takes a negative value when a hamlet lies on the spray-off side. If a hamlet is affected by multiple spray flights, we compute this number for each flight, and take the maximum of these numbers across all spray runs.¹⁷ A hamlet is then classified as treatment-side if this value is positive (i.e., spray-on side), and as control-side if it is negative (i.e., spray-off side).

The coding procedure involves several steps. First, we extend the recorded flight paths of fixed-wing aircraft missions from the flight start, turn, and end points by an arbitrary length of 0.05 degree \approx 5.56 km. Second, we add buffers of 0.001 degree \approx 111 meter width to both sides (222 meters) of the extended flight paths, as graphically illustrated by Figure 2(a). The extension length of 5.56 km and the 222 meter buffer width mimic ordinary spray flights. A routine fixed-wing aircraft mission involved multiple aircraft and dispersed herbicide at the airspeed of 130–150 knots (240–278 km/h) and an altitude of 150 feet, with each aircraft covering a laterally contiguous, partially overlapped swath of 73 to 80 meter (240 foot) wide and 16 km long (Buckingham, 1982, 37, 132; Darrow et al., 1969, 26; IoM, 1995, 25, 86–87; Stellman et al., 2003*b*, 327). An extension length of 5.56 km approximately

¹⁷Accordingly, if a hamlet is on the spray-off side for all flights, its running variable becomes the “minimum” absolute distance (i.e., the maximum of these negative numbers). Otherwise, it is the “maximum” distance among those missions that the hamlet location is on the spray-on side.

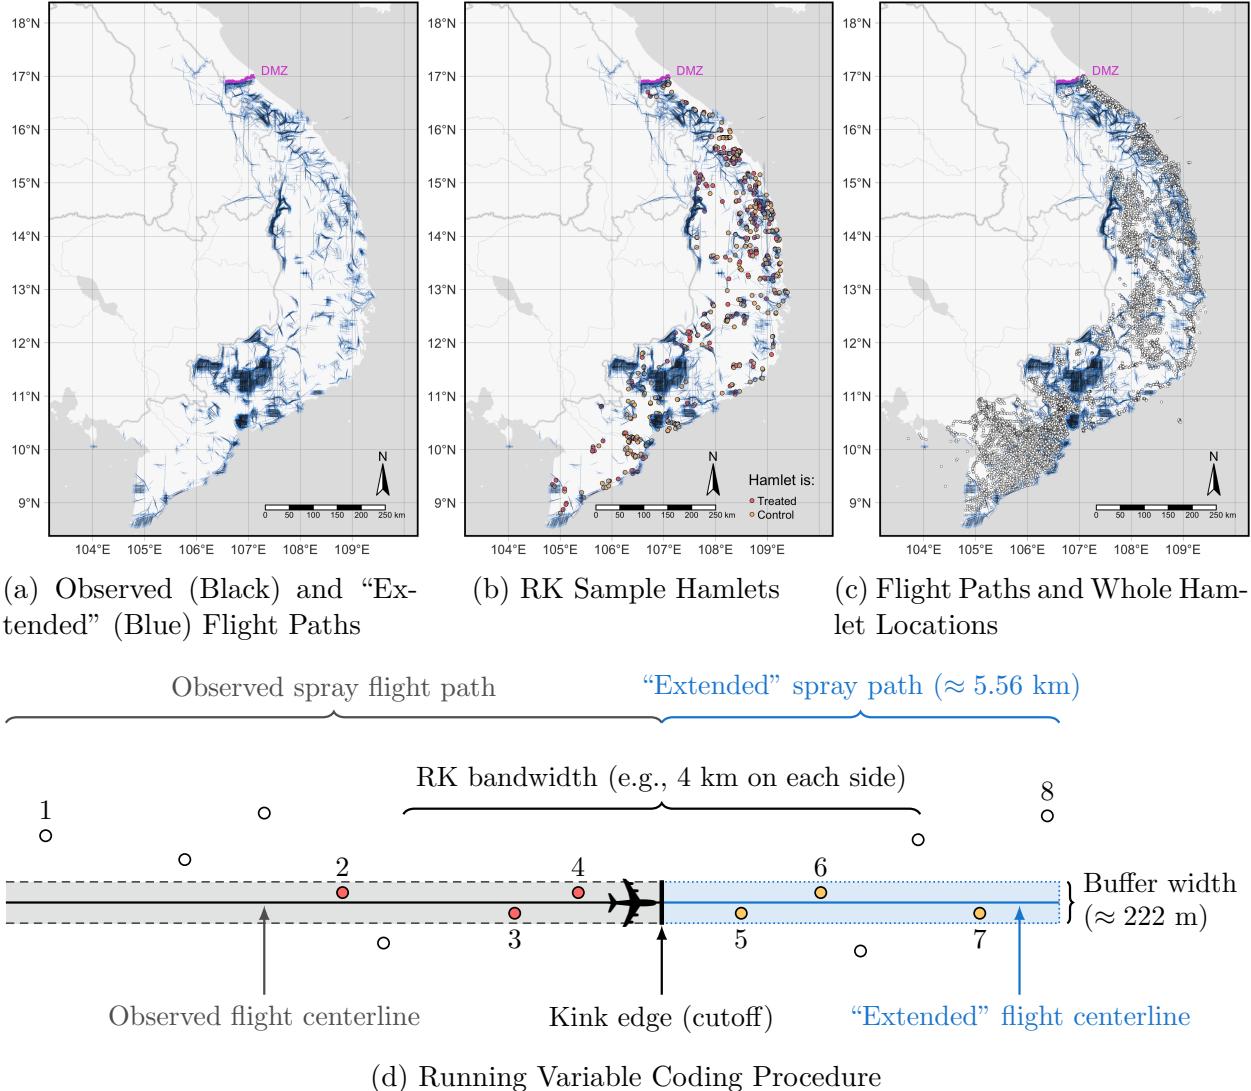


Figure 2: Spray Flight Paths, Hamlet Locations, and the RK Running Variable Coding Procedure
Notes: (a) Black polygons represent observed flight paths, and blue polygons indicate the extended paths with a 0.05 degree (≈ 5.56 km) distance from the start/turn/end points with a 111 meter buffer on each side (222 meters in total). Of the 14,733 hamlets (Panel (c)), hamlets covered by observed flight polygons with a 222 meter buffer (indicated by a gray polygon and black dashed segments) are categorized into the treatment groups (red dots), while the hamlets only covered by a blue polygon are categorized into the control group (orange dots; Panel (b)). (d) Vertical thick segment indicates a kink edge, and airplane symbol  indicates the direction of the spray flight. Dots and horizontal solid segments, respectively, represent hamlet locations and observed and “extended” spray flight paths. Dashed and dotted segments and shades represent the buffer around the observed (extended) spray path. The full RK sample keeps hamlets located within the 5.56 km extension length (e.g., hamlets 2–4 in the treatment group and hamlets 5–7 in the control group) while discarding the remaining hamlets (e.g., hamlets 1 and 8) located outside of the spray path polygon. A 4 km bandwidth further discards hamlets 2 and 7 from the analysis.

corresponds to the “less-than-one-half-minute-away” distance from the observed kink point at the airspeed of 240–278 km/h ($= 4\text{--}4.63 \text{ km/m}$). This extension length also corresponds

to typical intra-commune distances. Given the average area of contemporary communes of 29.6 km^2 , this spatial scale is equivalent to a square with a side length of approximately $5.44 \text{ km} (= \sqrt{29.6})$ or a circle with a diameter of roughly $6.14 \text{ km} (= 2 \times \sqrt{29.6/\pi})$. A 222 meter buffer width similarly approximates the combined swath of a routine spray mission involving three airplanes, with relatively limited lateral drifts (Darrow et al., 1969). This 222 meter or three-aircraft approximation also reflects the historical fact that three C-123 Provider airplanes were assigned to Operation Ranch Hand until 1964, and the number of C-123s increased to 12 in 1965 and then to 36 in 1967 (IoM, 1995, 86).

Third, we overlay the extended flight path polygons onto hamlet locations. The hamlets covered by the observed flight path are coded as “treated” (red dots), while those only covered by the extended parts are categorized into the control group (orange dots in Figures 2(c)). For example, in Figure 2(d), among the hamlets located within the 5.56 km distance from the kink edge (cutoff) in the middle, hamlets 2–4 are categorized into the treatment group while hamlets 5–8 into the control group. We then measure the geodesic distance between hamlet locations and the nearest kink edges to assign the running variable.

The geoprocessing leaves us a sample of 716 hamlets (4.86% of 14,733 hamlets in panel (c)) in 38 provinces in South Vietnam.¹⁸ As in Figure 2(b), the full RK sample includes the hamlets that are geographically covered by the observed or extended flight path polygons with the kink edge distance smaller than 5.56 km extension length.

¹⁸See Appendix A.3 for details of the running variable coding procedure including the multiple-path cases. Note also that the majority of hamlets were affected by three or fewer flight paths. In the baseline RK subsample with a 4 km bandwidth, this accounts for 467 out of 519 hamlets (89.98%). Specifically, 369 hamlets were covered by a single path (127 treated and 242 control), 84 hamlets by two paths (17 treated and 67 control), and 14 hamlets by three paths (4 treated and 10 control). The results reported below remain qualitatively unchanged with the limited subsample with 467 hamlets.

4.2 Model Specification

The following analysis builds upon two classes of estimation models. First, the distance-based, reduced-form model is specified as:

$$Y_{hp} = \gamma \text{EdgeDist}_{hp} + \tau^{\text{ITT}} \underbrace{\text{EdgeDist}_{hp} \times \mathbb{1}[\text{EdgeDist}_{hp} \geq 0]}_{\text{=Slope Change}} + \zeta \mathbb{1}[\text{EdgeDist}_{hp} \geq 0] + \mathbf{X}_{hp}^\top \boldsymbol{\beta} + \eta_p^{\text{Province}} + \kappa_{a[h]}^{\text{Agent}} + \theta_{l[h]}^{\text{Leg-end}} + \iota_{l[h]}^{\text{Phase}} + f(\text{Lon}_{hp}, \text{Lat}_{hp}) + e_{hp}, \quad (2)$$

where h , p , a , and l , respectively, index hamlets, provinces, herbicide agents, and flight leg edges. EdgeDist_{hp} is the geodesic distance from the kink edges in kilometers recentered at zero, and $\mathbb{1}[\text{EdgeDist}_{hp} \geq 0]$ is a dummy variable which takes one if $\text{EdgeDist}_{hp} \geq 0$ (treatment group hamlets) and zero otherwise (control group hamlets).¹⁹ \mathbf{X}_{hp} is a vector of hamlet-level covariates, and $f(\text{Lon}_{hp}, \text{Lat}_{hp})$ is a two-dimensional cubic polynomial of hamlet geocoordinates to screen out spatial trends.²⁰ η_p^{Province} denotes South Vietnam province fixed effects, and $\kappa_{a[h]}^{\text{Agent}}$ herbicide agent fixed effects indicating the chemical contamination with dioxin.²¹ $\theta_{l[h]}^{\text{Leg-end}}$ is leg-end fixed effects which take one if a hamlet's running variable is its distance to a spray leg end edge and zero otherwise; and $\iota_{l[h]}^{\text{Phase}}$ is fixed effects indicating a hamlet's running variable measured as the distance to the flight leg edges during three mission phases, 1962–1964, 1965–1966, and 1967–1971, reflecting the increase of the number of aircraft from 3 to 12 in 1965 and then to 36 in 1967 assigned to Operation Ranch Hand.²²

Second, to facilitate intuitive interpretation of the coefficient estimates, our RK estimation plugs-in the EOI of Equation (1) to the reduced-form model. Specifically, with abuse

¹⁹The specification includes $\mathbb{1}[\text{EdgeDist}_{hp} \geq 0]$ and thereby allows discontinuity in the treatment function at the kink point. When the kernel is symmetric (e.g., a uniform kernel), the asymptotic bias and variance of the RK estimand are not affected by the continuity imposition (Card et al., 2012, 2015b).

²⁰Linear, quadratic, and fifth-order polynomials of longitude and latitude yield qualitatively similar results.

²¹The agent fixed effects have two categories to reflect the chemical contamination with dioxin: TCDD (“Orange,” “Pink,” and “Purple”) and non-TCDD (“Blue” and “White”). We put eight records with an “Unknown” agent into the non-TCDD category.

²²As noted in Section 3.2, a “leg” indicates a part or a total path of a spray path represented as a straight line without turns. Leg-end fixed effects take the value of 1 for both the end and direction-change points.

of notations, the RK estimation builds upon the following two-stage model:

$$\underbrace{\ln \text{HERB}_{hp}}_{=\text{EOI}} = \gamma \text{EdgeDist}_{hp} + \delta \underbrace{\text{EdgeDist}_{hp} \times \mathbb{1}[\text{EdgeDist}_{hp} \geq 0]}_{=\text{Instrumental Variable}} + \zeta \mathbb{1}[\text{EdgeDist}_{hp} \geq 0] + \mathbf{X}_{hp}^\top \boldsymbol{\beta} + \eta_p^{\text{Province}} + \kappa_{a[h]}^{\text{Agent}} + \theta_{l[h]}^{\text{Leg-end}} + \iota_{l[h]}^{\text{Phase}} + f_1(\text{Lon}_{hp}, \text{Lat}_{hp}) + e_{hp}, \quad (3)$$

$$Y_{hp} = \tau^{\text{RK}} \underbrace{\ln \widehat{\text{HERB}}_{hp}}_{=\text{Predicted EOI}} + \lambda \text{EdgeDist}_{hp} + \nu \mathbb{1}[\text{EdgeDist}_{hp} \geq 0] + \mathbf{X}_{hp}^\top \boldsymbol{\xi} + \pi_p^{\text{Province}} + \rho_{a[h]}^{\text{Agent}} + \phi_{l[h]}^{\text{Leg-end}} + \psi_{l[h]}^{\text{Phase}} + f_2(\text{Lon}_{hp}, \text{Lat}_{hp}) + u_{hp}, \text{ if } |\text{EdgeDist}_{hp}| \leq b, \quad (4)$$

where HERB denotes the EOI, and the remaining terms are defined analogously to the reduced-form specification. The panel specification replaces the second-stage outcome with the annual population growth rate and adds a year fixed effect to the right hand side.²³ Following Card et al. (2012, 2015b, 2017), our estimation relies on an uniform (rather than triangular) kernel, and the baseline setup uses a bandwidth of $b = 4$ kilometers (519 hamlet observations). Intuitively, a 4 km bandwidth approximates the “one-minute away” distance at the typical airspeed of 4–4.63 km/m (= 240–278 km/h) of fixed-wing aircraft missions. To ensure that the fixed bandwidth choice does not drive the results, we replicate the RK estimates with alternative bandwidth sizes.

The parameters of interest are τ^{ITT} in the reduced form (Equation 2) and $\tau^{\text{RK}} = \frac{\tau^{\text{ITT}}}{\delta}$ in the second stage (Equation 4). τ^{ITT} and τ^{RK} , respectively, capture the intention-to-treat (ITT) effect of the spray-on/off or direction switching of fixed-wind missions and the average effect of a marginal increase in $\ln \text{HERB}$ on the outcome at the cutoff. Formally, τ^{RK} can be interpreted as the treatment-on-the-treated (TT) effect (Florens et al., 2008) or the local average response (LAR) of the herbicide EOI (Altonji and Matzkin, 2005), instrumented by

²³With abuse of notations, the panel version of the second-stage model is:

$$Y_{hpt} = \tau^{\text{RK}} \ln \widehat{\text{HERB}}_{hp} + \lambda \text{EdgeDist}_{hp} + \nu \mathbb{1}[\text{EdgeDist}_{hp} \geq 0] + \mathbf{X}_{hp}^\top \boldsymbol{\xi} + \pi_p^{\text{Province}} + \rho_{a[h]}^{\text{Agent}} + \phi_{l[h]}^{\text{Leg-end}} + \psi_{l[h]}^{\text{Phase}} + \eta_t^{\text{Year}} + f_2(\text{Lon}_{hp}, \text{Lat}_{hp}) + e_{hpt},$$

where η_t^{Year} represents year fixed effect, with the corresponding first-stage model specified analogously.

$\text{EdgeDist}_{hp} \times \mathbb{1}[\text{EdgeDist}_{hp} \geq 0]$ in the first stage (Equation 3, Card et al., 2015*b*).²⁴

Several aspects of the model specification warrant explanations. Given the asymptotic results of Pei et al. (2022), we adopt a local linear specification given the relatively small sample size. We also rely on province, agent, leg-end, and mission-phase fixed effects instead of the ideal flight or leg fixed effects due to the limited sample size. Nonetheless, the specification reflects the historical fact that South Vietnam Government both at the national and provincial levels and the U.S. jointly controlled the herbicide missions (Buckingham, 1982, 36–38; IoM, 1995, 86).²⁵ Province fixed effects subsume the regional variation in the authorization process and other province-level differences of counterinsurgency strategies; and the remaining fixed effects absorb the types of chemical shocks (agent fixed effects), aircraft directions (leg-end fixed effects), and overall intensity of spray missions and potential military significance of nearby targets (mission-phase fixed effects).

Throughout the analysis, we rely on the ordinary least squares (OLS) and the two-stage least-squares (2SLS) estimators. To account for the increased error due to the two-stage estimation and potential error dependence across space, we report Conley's (1999) standard errors robust to spatial clustering with a 30 km cutoff.²⁶ Appendix A reports descriptive statistics of the variables.

4.3 Identification Assumption

A valid fuzzy RK design hinges on the smoothness assumption, which yields two testable implications: First, the density of the running variable is sufficiently smooth, or continuously

²⁴Seen from the instrumental variable (IV) perspective, a key identification assumption in a RK design is that the instrument (spray-on side indicator and the running variable interaction term) affects the outcome (slope change) only through the EOI (slope change). Violation of this exclusion restriction condition invalidates the smoothness assumption (Card et al., 2015*b*). We briefly validate the identification assumption below and relegate the details to Appendix B.

²⁵The authorization process also involved the U.S. Ambassador, the U.S. Military Assistance Command, and the Corps Tactical Zones (CTZs). Dell and Querubin (2018) use the CTZ (Corps I-II) boundary as one of the sources for causal identification. As the CTZ boundaries follow the province boundaries, province fixed effects subsume the difference in counterinsurgency strategy across CTZs and local governments.

²⁶We use a 30 km cutoff to reflect the combat range of aircraft missions in which a C-123 airplane covered a swath of 80 m wide and 16 km long (Buckingham, 1982, 132). A 30 km cutoff approximately reflects the 16-km range as a radius. The sample mean (median) of the aircraft spray legs is 18 km (16.48 km).

differentiable at the cutoff; and second, predetermined covariates evolve smoothly around the kink point (Card et al., 2012, 2015*a*, 2017).

Following literature (Bana et al., 2020; Card et al., 2015*a,b*; Landais, 2015), we empirically validate the smoothness assumption with the running variable and covariate distributions in three ways. First, Figure B.1 in the Appendix examines continuity of the running variable distribution using the polynomial estimator of Cattaneo et al. (2020), and fails to detect statistically significant discontinuity at the kink point ($t = -0.209$ and $p = 0.835$). Second, we test for the kink in the running variable distribution using polynomial regressions.²⁷ As reported in Figure B.2 in the Appendix, the polynomial regressions fail to detect a slope change in the running variable distribution around the kink point or deterministic sorting. Finally, Figure B.3 in the Appendix presents a series of placebo kink estimates with the covariates, *including* hamlet population in 1967–1969 (during Operation Ranch Hand), as the left-hand-side variable using a reduced-form specification (Equation 2). Again, and consistent with the smoothness assumption, most of the placebo regressions reveals little discernible slope change at the kink point. Exceptions are the substantively small but statistically significant slope change in the distance to U.S. bases (standardized coefficient = 0.078; $t = 2.559$). Nonetheless, as in Appendix B.2, the same randomization inference exercise introduced in the next section disagrees with the asymptotic inference results and fails to negate that the covariate kink arises by random chance. Moreover, Appendix C.1 finds little evidence of model dependence, suggesting that the covariate kink merely plays a minor role in influencing the main results.

²⁷Formally, and with abuse of notation, we first aggregate the hamlet observations using bins with different sizes based on the running variable and then estimate the following regression model:

$$N_b^{\text{obs.}} = \beta \mathbb{1}[\overline{\text{EdgeDist}}_b \geq 0] + \sum_{p=1}^P \left[\gamma_p \overline{\text{EdgeDist}}_b^p + \delta_p \overline{\text{EdgeDist}}_b^p \times \mathbb{1}[\overline{\text{EdgeDist}}_b \geq 0] \right] + e_b,$$

where b indexes bins, $N_b^{\text{obs.}}$ reflects the number of observations in each bin, $\overline{\text{EdgeDist}}$ is the midpoint of EdgeDist of each bin, and P is the polynomial order. δ_1 captures the slope change of the probability density function of the running variable at the kink point, which is expected to be indistinguishable from zero.

5 Results

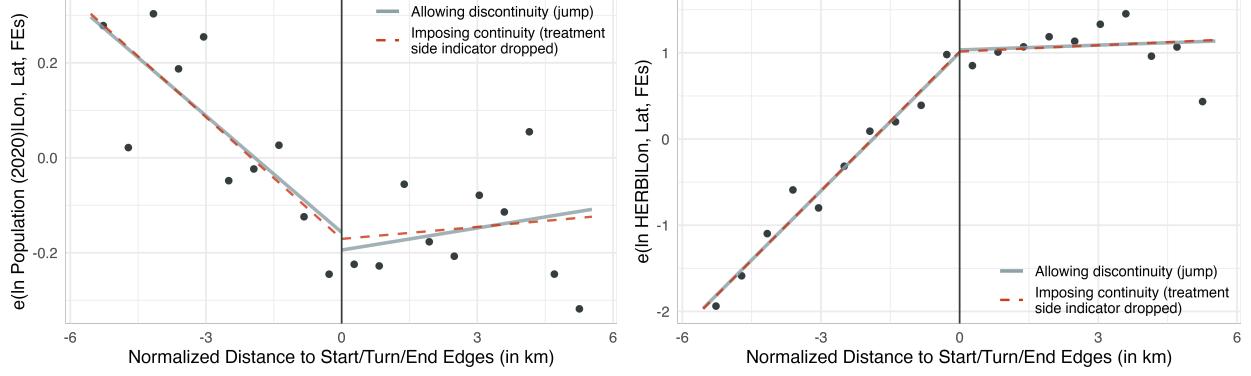
Figure 3 displays the distributions of population size in 2020 (reduced-form association; panel (a)) and the EOI (first-stage association; panel (b)) across the cutoff, given the two-dimensional polynomial of hamlet geocoordinates and the fixed effects. As explained above, our main estimation is the impact of herbicide spraying on population in terms of the distance from the spray on-off and aircraft-turning points. Figure 3(a) graphically illustrates the estimated distance gradient, or τ^{ITT} in the reduced-form specification (Equation 2), and Table C.1 in the Appendix reports the reduced-form estimates in regression format with and without additional covariate adjustments. The estimates show that population is markedly lower in the spray-on area (i.e., the right of the cutoff; Figure 3(a)). On the spray-off side (i.e., the left of the cutoff) population size gradually recovers as the distance increases from the cutoff, consistent with the possible drift of herbicide beyond the cutoff.

As noted earlier, we reinterpret this main result in terms of herbicide quantity. We do so by linking the distance to the EOI, which we call the “first stage,” as the index reflects the “built-in” slope induced by construction of Equation (1) rather than the “ground truth” levels of herbicide exposure. Figure 3(b) presents this first-stage association, or δ in Equation 3. Using the EOI as a “pseudo” treatment variable, we then reestimate the treatment effect, a step we refer to as the “second stage.”²⁸ In the following, we first report the cross-sectional RK estimates with population size as the outcome. We then extend the analysis with growth rate and panel specifications, followed by a brief summary of robustness checks.

²⁸The fuzzy RK estimand, τ^{RK} , can be written as the ratio of the reduced-form slope change (i.e., slope change in the outcome function, τ^{ITT}) relative to the first-stage slope change (i.e., slope change in the treatment function, δ) at the kink point as (Card et al., 2012, 2015b):

$$\tau^{\text{RK}} = \left(\frac{\lim_{v_0 \downarrow 0} d(\mathbb{E}[Y|V=v])}{dv} \Big|_{v=v_0} - \frac{\lim_{v_0 \uparrow 0} d(\mathbb{E}[Y|V=v])}{dv} \Big|_{v=v_0} \right) \Big/ \left(\frac{\lim_{v_0 \downarrow 0} d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0} - \frac{\lim_{v_0 \uparrow 0} d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0} \right),$$

where V denotes the running variable, v_0 the kink point, and D the treatment.



(a) Outcome Kink: $\ln \text{Population (2020)}$

(b) EOI Kink: $\ln \text{HERB}$

Figure 3: Discontinuous Slope Change in the Outcome and EOI Functions

Notes: Dots represent the binned frequencies, and solid and dashed lines show linear regression fits adjusting for longitude, latitude, and the fixed effects. Solid line allows discontinuity at the kink point (as in Equations 2 and 3), while the dashed line imposes continuity by dropping the treatment-side dummy variable. The bin size is selected by the integrated mean squared error-optimal mimicking variance evenly-spaced method using spacing estimators (Calonico et al., 2015).

5.1 Population Size Estimates

Plugging the EOI score into the reduced-form model, Table 1 reports the more intuitively interpretable, RK estimation results with population size in 2001 (Panel B) and 2020 (Panel C) as the outcome, along with the first-stage estimates (Panel A). Model (1) is the baseline model without additional controls besides the two-dimensional polynomial of longitude and latitude and fixed effects. Models (2) to (4) consecutively add the proximity to the key targets, geographical characteristics, and historical attributes as covariates. To address potential spillover effects, Model (5) adjusts for the spatially-lagged EOI with the 30 km neighbor cutoff in addition to the full set of covariates. Note that given the limitation in data availability and the historical facts, some of the covariates can partly be posttreatment (e.g., bombing point distance, 1965–1971). Thus, the models with covariate adjustments do not necessarily provide conservative estimates due to possible posttreatment bias of unknown directions and should be interpreted with caution.

Table 1 confirms the graphical evidence of Figure 3 and underlines the herbicide legacies. Across model specifications, the first-stage, EOI kink estimates are substantively and statis-

Table 1: Herbicide Exposure and Population Size in 2001 and 2020

	Panel A: ln HERB				
	(1)	(2)	(3)	(4)	(5)
First stage					
δ : EdgeDist \times $\mathbb{1}[\text{EdgeDist} \geq 0]$	-0.569*** (0.081)	-0.552*** (0.075)	-0.565*** (0.072)	-0.564*** (0.070)	-0.564*** (0.070)
<i>F</i> -statistic (weak instrument)	49.209	54.014	61.924	64.440	65.179
Adjusted R ²	0.570	0.609	0.613	0.617	0.617
Panel B: ln Population (2001)					
	(1)	(2)	(3)	(4)	(5)
Second stage					
τ^{RK} : ln HERB	-0.264*** (0.082)	-0.262*** (0.076)	-0.211*** (0.069)	-0.198*** (0.071)	-0.190*** (0.071)
10% increase effect size	-2.51%	-2.50%	-2.01%	-1.88%	-1.81%
Average outcome	1.311	1.311	1.311	1.311	1.311
Randomization inference (reduced form)					
Kink point (KP) randomization <i>p</i> -value	0.019	0.009	0.008	0.010	0.011
Running variable (RV) randomization <i>p</i> -value	0.003	0.001	0.002	0.005	0.007
Spatial noise (SN) randomization <i>p</i> -value	0.003	0.000	0.002	0.005	0.006
Panel C: ln Population (2020)					
	(1)	(2)	(3)	(4)	(5)
Second stage					
τ^{RK} : ln HERB	-0.343*** (0.088)	-0.346*** (0.082)	-0.284*** (0.075)	-0.274*** (0.078)	-0.266*** (0.078)
10% increase effect size	-3.27%	-3.29%	-2.70%	-2.61%	-2.53%
Average outcome	1.464	1.464	1.464	1.464	1.464
Randomization inference (reduced form)					
KP randomization <i>p</i> -value	0.022	0.012	0.010	0.012	0.013
RV randomization <i>p</i> -value	0.001	0.000	0.001	0.001	0.000
SN randomization <i>p</i> -value	0.000	0.000	0.000	0.001	0.001
Observations	519	519	519	519	519
Avg. N neighbors (Conley SE cluster size)	27.5	27.5	27.5	27.5	27.5
Key target covariates	✓	✓	✓	✓	✓
Geographic covariates		✓	✓	✓	✓
Historical covariates			✓	✓	✓
ln spatially-lagged HERB					✓
Fixed effects and $f(\text{Lon}, \text{Lat})$	✓	✓	✓	✓	✓

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Conley (1999) standard errors adjusted for spatial clustering with a 30 km cutoff and a Bartlett kernel are in parentheses. **Key target covariates**: NVA base distance, population (1967–1969), U.S. base distance, U.S. troop distance, forest presence, rice cropland, road distance, slash and burn cropland, Viet Cong control prevalence. **Geographic covariates**: Precipitation, temperature, u - and v -components of neutral wind, elevation, flow accumulation, rice suitability, river distance, ruggedness. **Historical covariates**: Bombing point distance, border distance, number of neighbor hamlets, railway distance. ln spatially-lagged HERB is the logged average HERB among the neighbor hamlets with a 30 km cutoff. **Fixed effects**: Agent fixed effect, province fixed effect, end-leg fixed effect, mission phase fixed effect. **Randomization inference**: KP randomization *p*-value is computed by the permutation test of Ganong and Jäger (2018) with 10,000 placebo kink points drawn from uniform distribution $\mathcal{U}(-4 \text{ km}, 4 \text{ km})$. RV randomization *p*-value is obtained by randomly assigning the running variable to the sample hamlets for 10,000 times. SN randomization *p*-value is obtained by 10,000 synthetic spatial noise simulations of Kelly (2021) with the outcome replaced by randomly generated spatial noise.

tically significant (Panel A). The second-stage coefficients on $\ln \text{HERB}$ are also consistently signed negative and retain the statistical significance at the conventional 5% level (Panels B and C). Note also that the uninstrumented OLS estimates in Appendix Table A.2 underestimate or even fail to reveal the negative association, suggesting that nonrandom herbicide assignment invites bias into naive comparisons. Moreover, the coefficient stability suggests that the estimates are unlikely to be driven by an arbitrary covariate adjustment choices, which is of a particular concern in RK applications (Ando, 2017) and we further investigate in Appendix C.1. The negative association also remains visible across the outcomes, population size in 2001 and 2020.

As an initial robustness check, Table 1 reports three randomization inference results.²⁹ First, an important concern for RK applications is the misspecification of the nonlinearity such that one falsely specifies a curvature without kink as a discontinuous slope change. To guard against the misspecification bias, we report kink point (KP) randomization p -value obtained from the reduced-form permutation test (Ganong and Jäger, 2018).³⁰ For the permutation exercise, we first generate 10,000 placebo kink points randomly drawn from a uniform distribution, $\mathcal{U}(-4 \text{ km}, 4 \text{ km})$. We then consecutively estimate the reduced-form version of the RK model with the subsample within baseline 4 km bandwidth around each placebo kink point to obtain the empirical distribution of placebo estimates. The two-sided p -value is computed by doubling the minimum of the fraction of placebo estimates no smaller than or no greater than the actual reduced-form coefficient. Second, another intuitive approach of randomization inference is to randomly assign the running variable (RV) and the corresponding treatment-side indicator variable to the sample hamlets while holding the kink point fixed (e.g., Dell and Querubin, 2018, 30–33; see also, Cattaneo et al., 2015). The RV randomization two-sided p -values report the share of the absolute 10,000 placebo

²⁹Randomization exercise relies on the reduced-form and first-stage models as randomization inference cannot be applied without additional assumptions to unobservable subgroups such as compilers.

³⁰If our RK specification correctly captures a discontinuous slope change, we would be unlikely to see a slope change with placebo kink points, and thus have a small KP randomization p -value. If, on the other hand, we falsely specify a quadratic function with no kink as a discontinuous slope change, we would observe discernible kink estimates with placebo kink points and a large KP randomization p -value.

coefficients that are larger than the absolute actual reduced-form kink coefficient. Finally, to address potential spatial curve-fitting and residual spatial autocorrelation, we use the spatial noise (SN) randomization inference developed by [Kelly \(2021\)](#). This procedure replaces the left-hand-side variable by synthetic noise with the same estimated spatial structure of the observed outcome partialled out by covariates. If the estimation is not an artifact of spatial trends, the observed variable should not explain spatial noise. The SN randomization p -value indicates the fraction of 10,000 spatial noise simulations which yield absolute t values greater than the absolute t value estimated with the observed data.

For both population size in 2001 and 2020, the randomization inference results are consistent with asymptotic inference. The three series of randomization p -values remain smaller than the conventional 5% threshold, suggesting that the actual estimates are unlikely to arise from misspecification of the underlying nonlinearity (KP randomization), random chance (RV randomization), or spatial autocorrelation and curve-fitting (SN randomization).

Turning to magnitude, the coefficient estimates can readily be interpreted as elasticity given the log-log specification. The estimates suggest that, measured at the 100 meter grid scale, a 10% increase in herbicide exposure is persistently followed by a 2.61% ($-0.274 \times \ln(1.1) = -0.0261$, Model 4, with all covariates) to 3.27% (Model 1, without covariates) decrease in population size in 2020 (Panel C), and a 1.88% to 2.51% decrease in 2001 (Models 1 and 4 in Panel B). Note that the 10% increase reference does not overstate the magnitude of the treatment effect, as such an increase in herbicide exposure is observed within a geographically small area around the kink edges. Specifically, a 10% increase in the treatment approximately corresponds to a 10.8% increase in $\ln \text{HERB}$ between the average of the control hamlets within a 250 meter distance from the kink edge and the average of the treated hamlets within the same distance from the kink edge.

Besides the baseline population growth during the 2001–2020 period (from 1.311 to 1.464 in the logarithm scale, Table 1), the growing, rather than fading away, coefficient size implies that herbicide legacies shape not only population *size* but also population *growth rate*.

This leads to a testable implication that a similar negative association is also present in local population growth rate, which contrasts the conventional wisdom of rapid recovery of population from wartime destruction.

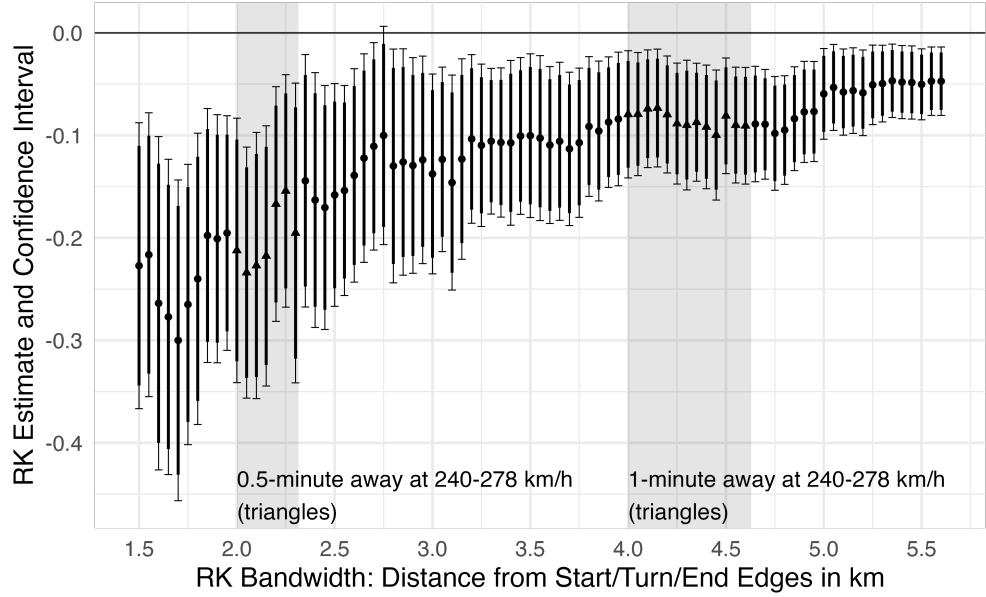
5.2 Population Growth Rate Estimates

To examine the testable implication of cumulatively increasing legacies of herbicide exposure, we reestimate the RK models with the overall and annual population growth rates during the 2001–2020 period as the outcomes, using both cross-sectional and panel setups. This growth-rate specification also surmounts the robustness concern that a small number of volatile and less persistent population count observations in 2001 and 2020 drive the estimation results. For the panel specification, we extend the baseline cross-sectional specification by replacing the outcome by the annual growth rate and adding a year fixed effect as an additional control.³¹ In the panel setup, we report robust standard errors clustered at the hamlet level. Another robustness issue remaining in Table 1 is that the results are based on a fixed bandwidth of 4 km, leaving a potential concern for the bandwidth choice driving the results. To examine the herbicide legacies on population growth rate and address these robustness concerns, we report the RK estimates for overall and annual population growth rates across different RK bandwidths ranging from 1.5 km to 5.6 km with increments of 50 meters.

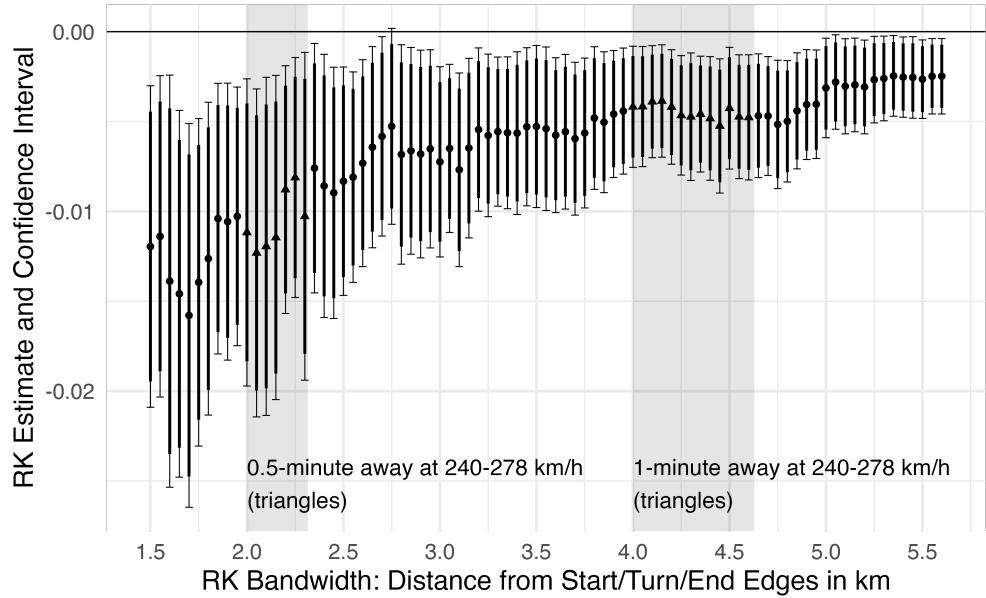
Figure 4 presents the RK estimates for overall (Panel (a)) and annual (Panel (b)) population growth rates across bandwidth sizes. Regardless of the bandwidth settings and outcomes, the growth rate estimates reveal a negative association consistent with the population size estimates.³² The effect size is also discernible given the average outcome of 0.153 (or $e^{0.153} = 1.165$; 16.5% overall population growth in 2001–2020) and the geographically small

³¹The results remain qualitatively unchanged when further adjusting for the lagged outcome, Y_{hdt-1} .

³²Figure 4 follows the specification of model (1) in Table 1. Note that, as reported in Appendix C.1, the growth rate specification is remarkably robust to different model specifications (different combinations of adjusted covariates) and exhibits little model dependence. The empirical distribution of the ITT estimates of all possible 2^N covariates = $2^{21} = 2,097,152$ model specifications, jointly including or excluding the u - and v -components of neutral winds, is almost normally distributed as in Appendix Figure C.1, suggesting that adjusted covariates affecting the ITT estimates only due to random error.



(a) Overall Population Growth Rate, 2001–2020



(b) Annual Population Growth Rate, 2002–2020

Figure 4: Herbicide Exposure and Population Growth Rate

Notes: Symbols and thin (thick) vertical segments with horizontal ticks represent the RK estimates and the corresponding 95% (90%) confidence intervals based on Conley (1999) standard errors with spatial clustering (Panel (a)) and robust standard errors with hamlet-level clustering (Panel (b)). The model specification in Panel (a) follows model (1) in Table 1. Panel (b) adds a year fixed effect to the baseline specification. Grey shades represent “0.5-minute-away” (2–2.315 km) and “one-minute-away” distances (4–4.63 km) from the kink point at the typical airspeed of Ranch Hand aircraft of 240–278 km/h, and triangles indicate the corresponding point estimates. Annual growth estimate in Panel (b) covers the 2002–2020 period as the outcome is defined as $\ln \text{Population}_t - \ln \text{Population}_{t-1}$ and our dataset covers the 2001–2020 period.

level of measurement of 100 meter grid. The cross-sectional RK estimate of $\tau = -0.079$ with the baseline bandwidth of 4 km suggests that a 10% increase in herbicide exposure translates into a 0.75 percentage point decrease ($-0.079 \times \ln(1.1) = -0.0075$) in the overall population growth rate in 2001–2020 (Figure 4(a)).³³ The negative association is also visible in the panel specification with the annual population growth rate as the outcome (Figure 4(b)), revealing the persistent legacies in the dynamics as well as the snapshots of population outcomes in the present day. The two panels of Figure 4 also indicate that the estimates are qualitatively stable across bandwidth choices and that the baseline bandwidth of 4 km provides conservative estimates compared to those obtained with smaller bandwidths.

5.3 Robustness Check and Sensitivity Analysis

Robustness and sensitivity concerns in the main estimates include (1) the robustness to the parameter choices to construct the RK sample and the exposure score, HERB, (2) confounding kink in observed covariates and model dependence, and (3) unobserved confounding kink. In addition to the placebo kink (covariate kink) estimates reported above and Appendix B.2, Appendix C addresses these concerns by leveraging alternative combinations of flight buffer width and half-decay distance parameters, and all possible 2^N covariates $= 2^{21} = 2,097,152$ model specifications per outcome, treating the u - and v -components of neutral winds as a single block.³⁴ We also report a jackknife approach and the known historical differences in counterinsurgency strategies across four U.S. Corps Tactical Zones (CTZs) in South Vietnam (see, e.g., Dell and Querubin, 2018), and a sensitivity analysis approach to quantify how severe unobserved confounding forces would need to be in order to eliminate the main estimates (Cinelli and Hazlett, 2020).³⁵ Reassuringly, none of the robustness checks and

³³Given $\ln \text{Population}_{2020} - \ln \text{Population}_{2001} \approx \frac{\text{Population}_{2020}}{\text{Population}_{2001}} - 1$, a 0.75% decrease can approximately be interpreted as the effect size in the percentage point scale.

³⁴To improve computational efficiency, we always jointly include or exclude the u - and v -components of neutral winds, while treating all remaining covariates independently.

³⁵During the Vietnam War, the U.S. divided South Vietnam into four CTZs commanded by different branches of the U.S. military. Dell and Querubin (2018) exploit the CTZ I–II boundary as one of the sources for causal identification, given the historical fact that the U.S. Marine Corps (USMC) in Corps Region I

sensitivity analysis yields results that would invalidate or overturn the main results, lending additional credibility to the findings.

6 Conclusion

This article combined the historical records of herbicide missions with a fuzzy RK design and revealed the negative legacies of herbicidal warfare on contemporary population size and growth rate in Vietnam. The empirical analysis suggests that the temporary shock of herbicidal warfare left lasting, rather than temporary, effects on contemporary population outcomes. While the reported results tell us little about the underlying mechanisms, there are at least three distinct, although not mutually exclusive, explanations for the revealed persistent associations, which also opens up pathways for future research. First, the negative association might reflect decreased birthrates and life expectancy. The reported negative association between Agent Orange exposure and health outcomes (e.g., [Le et al., 2022](#); [Vuong, 2024](#); [Yamashita and Trinh, 2022](#)) makes the heavily sprayed areas persistently suffer lower birthrates and decreased life expectancy. Second, domestic migration preferences and patterns, beyond the temporary displacement of peasants to the city during the Vietnam War ([Huntington, 1968](#)), can also contribute to population dynamics. Increased out-migration from the herbicide affected areas coupled with decreased in-migration into the damaged areas can also account for the lasting negative associations. Finally, the lower population size and growth rates can arise from deteriorated agricultural productivity to sustain local population (e.g., [Appau et al., 2021](#)).

While the limitation in the available hamlet-level records prevents focused analysis, the first mechanism aligns with the insights of the existing epidemiological and economic research on the exposure to pesticide. As noted earlier, for example, [Camacho and Mejía \(2017\)](#), [Calzada et al. \(2023\)](#); [Dias et al. \(2023\)](#), and [Frank \(2024\)](#) highlight the increased

implemented a moderate hearts-and-minds-oriented strategy while the U.S. Army commanded Corps Region II and emphasized overwhelming firepower and large-scale operations.

infant mortality in areas with past exposure and usage of pesticide. While the findings are based on the relatively short-termed records, worsening birthrates coupled with deteriorated health outcomes would, at least partly, explain the decreased population outcomes in the longer term, as documented in this article (see also, Le et al., 2022; Vuong, 2024; Yamashita and Trinh, 2022). Ultimately, our findings suggest that herbicidal warfare, along with carpet-bombing, during the Vietnam War induced more than just temporary displacement or “forced-draft urbanization and modernization” (Huntington, 1968, 652); instead, it served as a catalyst for a lasting demographic erosion that continues to suppress both population size and growth trajectories.

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process

During the preparation of this article, the authors used ChatGPT-5 and Gemini 3 to improve the English phrasing and grammar. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the article.

References

Altonji, Joseph G., and Rosa L. Matzkin. 2005. “Cross section and panel data estimators for nonseparable models with endogenous regressors.” *Econometrica*, 73(4): 1053–1102.

Ando, Michihito. 2017. “How much should we trust regression-kink-design estimates?” *Empirical Economics*, 53(3): 1287–1322.

Appau, S, S A Churchill, R Smyth, and T A Trinh. 2021. “The long-term impact of the Vietnam War on agricultural productivity.” *World Development*, 146(October): 105613.

Bana, Sarah H., Kelly Bedard, and Maya Rossin-Slater. 2020. “The Impacts of Paid Family Leave Benefits: Regression Kink Evidence from California Administrative Data.” *Journal of Policy Analysis and Management*, 39(4): 888–929.

Bauer, Michal, Christopher Blattman, Julie Chytilová, Joseph Henrich, Edward Miguel, and Tamar Mitts. 2016. “Can War Foster Cooperation?” *Journal of Economic Perspectives*, 30(3): 249–274.

Berman, Chantal, Killian Clarke, and Rima Majed. 2023. “From Victims to Dissidents: Legacies of Violence and Popular Mobilization in Iraq (2003-2018).” *American Political Science Review*, forthcoming.

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

Bosker, Maarten, Steven Brakman, Harry Garretsen, and Marc Schramm. 2007. “Looking for multiple equilibria when geography matters: German city growth and the WWII shock.” *Journal of Urban Economics*, 61(1): 152–169.

Brainerd, Elizabeth, and Nidhiya Menon. 2014. “Seasonal Effects of Water Quality: The Hidden Costs of the Green Revolution to Infant and Child Health in India.” *Journal of Development Economics*, 107: 49–64.

Brakman, Steven, Harry Garretsen, and Marc Schramm. 2004. “The strategic bombing of German cities during World War II and its impact on city growth.” *Journal of Economic Geography*, 4(2): 201–218.

Buckingham, William A. 1982. *Operation Ranch Hand: The Air Force and Herbicides in Southeast Asia, 1961-1971*. Washington, DC: Office of Air Force History.

Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik. 2015. “Optimal Data-Driven Regression Discontinuity Plots.” *Journal of the American Statistical Association*, 110(512): 1753–1769.

Calzada, Joan, Meritxell Gisbert, and Bernard Moscoso. 2023. “The Hidden Cost of Bananas: The Effects of Pesticides on Newborns’ Health.” *Journal of the Association of Environmental and Resource Economists*, 10(6): 1623–1663.

Camacho, Adriana, and Daniel Mejía. 2017. “The Health Consequences of Aerial Spraying Illicit Crops: The Case of Colombia.” *Journal of Health Economics*, 54: 147–160.

Card, David, Andrew Johnston, Pauline Leung, Alexandre Mas, and Zhuan Pei. 2015a. “The effect of unemployment benefits on the duration of unemployment insurance receipt: New evidence from a regression kink design in Missouri, 2003–2013.” *American Economic Review*, 105(5): 126–130.

Card, David, David Lee, Zhuan Pei, and Andrea Weber. 2012. “Nonlinear Policy Rules and the Identification and Estimation of Causal Effects in a Generalized Regression Kink Design.”

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber. 2015b. “Inference on Causal Effects in a Generalized Regression Kink Design.” *Econometrica*, 83(6): 2453–2483.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber. 2017. “Regression Kink Design: Theory and Practice.” In *Regression Discontinuity Designs: Theory and Applications*. , ed. Matias D. Cattaneo and Juan Carlos Escanciano. Bingley: Emerald Publishing.

Cattaneo, Matias D., Brigham R. Frandsen, and Rocío Titiunik. 2015. “Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate.” *Journal of Causal Inference*, 3(1): 1–24.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma. 2020. “Simple Local Polynomial Density Estimators.” *Journal of the American Statistical Association*, 115(531): 1449–1455.

Cinelli, Carlos, and Chad Hazlett. 2020. “Making sense of sensitivity: Extending omitted variable bias.” *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 82(1): 39–67.

Conley, Timothy G. 1999. “GMM estimation with cross sectional dependence.” *Journal of Econometrics*, 92(1): 1–45.

Consonni, D., A. C. Pesatori, C. Zocchetti, R. Sindaco, L. C. D’Oro, M. Rubagotti, and P. A. Bertazzi. 2008. “Mortality in a Population Exposed to Dioxin after the Seveso, Italy, Accident in 1976: 25 Years of Follow-Up.” *American Journal of Epidemiology*, 167(7): 847–858.

Darrow, Robert A., Kent R. Irish, and Charles E. Minarik. 1969. “Herbicides Used in Southeast Asia.” *USDA National Agricultural Library*. Accessed November 16, 2025, <https://www.nal.usda.gov/exhibits/speccoll/items/show/2304>.

Davis, Donald R., and David E. Weinstein. 2002. “Bones, Bombs, and Break Points: The Geography of Economic Activity.” *American Economic Review*, 92(5): 1269–1289.

Davis, Donald R., and David E. Weinstein. 2008. “A search for multiple equilibria in urban industrial structure.” *Journal of Regional Science*, 48(1): 29–65.

De Araujo, Leandro Goulart, Daniel Froes Zordan, Alain Celzard, and Vanessa Fierro. 2023. “Glyphosate Uses, Adverse Effects and Alternatives: Focus on the Current Scenario in Brazil.” *Environmental Geochemistry and Health*, 45(12): 9559–9582.

Defense Digital Service. 2016. “Theater History of Operations Data (THOR): Vietnam War.” Available at: <https://data.world/datamil/vietnam-war-thor-data>.

Dell, Melissa, and Pablo Querubin. 2018. “Nation Building Through Foreign Intervention: Evidence from Discontinuities in Military Strategies.” *Quarterly Journal of Economics*, 133(2): 701–764.

Dias, Mateus, Rudi Rocha, and Rodrigo R. Soares. 2023. “Down the River: Glyphosate Use in Agriculture and Birth Outcomes of Surrounding Populations.” *Review of Economic Studies*, 90(6): 2943–2981.

Douglass, Rex W. 2011. “Data of the Vietnam War v1.0 (Vietnam Hamlet Evaluation System Gazetteer Data).” Available at: <https://esoc.princeton.edu/data/vietnam-hamlet-evaluation-system-gazetteer-data>, accessed July 7, 2021.

Dung, Phan Xuan. 2023. “Agent Orange Victims in Vietnam: Their Numbers, Experiences, Needs, and Sources of Support.” United States Institute of Peace, 2301 Constitution Avenue NW Washington, DC 20037.

Fick, Stephen E., and Robert J. Hijmans. 2017. “WorldClim 2: New 1-km spatial resolution climate surfaces for global land areas.” *International Journal of Climatology*, 37(12): 4302–4315.

Florens, J.P., J.J. Heckman, C Meghir, and E. Vytlacil. 2008. “Identification of Treatment Effects Using Control Functions in Models With Continuous, Endogenous Treatment and Heterogeneous Effects.” *Econometrica*, 76(5): 1191–1206.

Frank, Eyal G. 2024. “The economic impacts of ecosystem disruptions: Costs from substituting biological pest control.” *Science*, 385(6713).

Ganong, Peter, and Simon Jäger. 2018. “A Permutation Test for the Regression Kink Design.” *Journal of the American Statistical Association*, 113(522): 494–504.

Grasse, Donald. 2023. “State Terror and Long-Run Development: The Persistence of the Khmer Rouge.” *American Political Science Review*, forthcoming.

Hersbach, Hans, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dianand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla, Xavier Abellán, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot, Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger, Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux, Philippe Lopez, Cristina Lupu, Gabor Radnóti, Patricia de Rosnay, Iryna Rozum, Freja Vamborg, Sébastien Villaume, and Jean Noël Thépaut. 2020. “The ERA5 global reanalysis.” *Quarterly Journal of the Royal Meteorological Society*, 146(730): 1999–2049.

Huntington, Samuel P. 1968. “The Bases of Accommodation.” *Foreign Affairs*, 46(4): 642.

Institute of Medicine (Committee to Review the Health Effects in Vietnam Veterans of Exposure to Herbicides). 1995. *Veterans and Agent Orange: Health Effects of Herbicides Used in Vietnam*. Washington, D.C.: National Academy Press.

International Agency for Research on Cancer (IARC). 2012. *Chemical Agents and Related Occupations: IARC Monographs on the Evaluation of Carcinogenic Risks to Humans, Volume 100F*. Lyon: IARC.

Kelly, Morgan. 2021. “Persistence, Randomization, and Spatial Noise.” *CEPR Discussion Paper*, DP16609.

Lai, Wangyang. 2017. “Pesticide Use and Health Outcomes: Evidence from Agricultural

Water Pollution in China." *Journal of Environmental Economics and Management*, 86: 93–120.

Landais, Camille. 2015. "Assessing the welfare effects of unemployment benefits using the regression kink design." *American Economic Journal: Economic Policy*, 7(4): 243–278.

Le, Duong Trung, Thanh Minh Pham, and Solomon Polacheck. 2022. "The long-term health impact of Agent Orange: Evidence from the Vietnam War." *World Development*, 155: 105813.

Lehner, Bernhard, Kristine Verdin, and Andy Jarvis. 2008. "New global hydrography derived from spaceborne elevation data." *Eos*, 89(10): 93–94.

Lichter, Andreas, Max Löffler, and Sebastian Siegloch. 2020. "The Long-Term Costs of Government Surveillance: Insights from Stasi Spying in East Germany." *Journal of the European Economic Association*, 19(2): 741–789.

Lin, Erin. 2022. "How War Changes Land: Soil Fertility, Unexploded Bombs, and the Underdevelopment of Cambodia." *American Journal of Political Science*, 66(1): 222–237.

McCormick, John T. 2021. "The Hamlet Evaluation System Reevaluated." Available at: https://github.com/jtmccorm/HES_Reevaluated, accessed July 7, 2021.

Miguel, Edward, and Gérard Roland. 2011. "The long-run impact of bombing Vietnam." *Journal of Development Economics*, 96(1): 1–15.

Military Assistance Command, Vietnam (MACV). 1969. "Military Operations: Herbicide Operations (U)." *APO San Francisco 96222, Accession Number: AD0779793*. Available at: <https://apps.dtic.mil/sti/citations/AD0779793>, accessed July 10, 2023.

Mocarelli, Paolo, Pier Mario Gerthoux, Enrica Ferrari, Donald G Patterson, Stephanie M Kieszak, Paolo Brambilla, Nicoletta Vincoli, Stefano Signorini, Pierluigi Tramacere, Vittorio Carreri, Eric J Sampson, Wayman E Turner, and Larry L Needham. 2000. "Paternal Concentrations of Dioxin and Sex Ratio of Offspring." *The Lancet*, 355(9218): 1858–1863.

National Academy of Sciences. 1974. *The Effects of Herbicides in South Vietnam: Part A Summary and Conclusions*. Washington, DC: National Academy of Sciences.

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

Palmer, Michael, Cuong Viet Nguyen, Sophie Mitra, Daniel Mont, and Nora Ellen Groce. 2019. "Long-lasting consequences of war on disability." *Journal of Peace Research*, 56(6): 860–875.

Pei, Zhuan, David S. Lee, David Card, and Andrea Weber. 2022. "Local Polynomial Order in Regression Discontinuity Designs." *Journal of Business and Economic Statistics*, 40(3): 1259–1267.

Riaño, Juan Felipe, and Felipe Valencia Caicedo. 2024. "Collateral Damage: The Legacy of the Secret War in Laos." *Economic Journal*, forthcoming.

Rowley, Ralph A. (Office of Air Force History). 1975. "The Air Force in Southeast Asia: FAC Operations, 1965–1970." *Office of Air Force History, May 1975*, available at: <https://media.defense.gov/2011/Mar/24/2001330117/-1/-1/0/AFD-110324-008.pdf>.

Singhal, Saurabh. 2019. "Early life shocks and mental health: The long-term effect of war in Vietnam." *Journal of Development Economics*, 141(March): 102244.

Skidmore, Marin Elisabeth, Kaitlyn M. Sims, and Holly K. Gibbs. 2023. "Agricultural Intensification and Childhood Cancer in Brazil." *Proceedings of the National Academy of Sciences*, 120(26): 10324–10333.

of Sciences, 120(45): e2306003120.

Stellman, Jeanne Mager, and Steven D. Stellman. 2018. “Agent orange during the Vietnam war: The lingering issue of its civilian and military health impact.” *American Journal of Public Health*, 108(6): 726–728.

Stellman, Jeanne Mager, Steven D. Stellman, Richard Christian, Tracy Weber, and Carrie Tomasallo. 2003a. “The extent and patterns of usage of agent orange and other herbicides in Vietnam.” *Nature*, 422(6933): 681–687.

Stellman, Jeanne Mager, Steven D. Stellman, Tracy Weber, Carrie Tomasallo, Andrew B. Stellman, and Richard Christian. 2003b. “A geographic information system for characterizing exposure to Agent Orange and other herbicides in Vietnam.” *Environmental Health Perspectives*, 111(3): 321–328.

Stellman, Steven D., and Jeanne M. Stellman. 1986. “Estimation of exposure to agent orange and other defoliants among american troops in vietnam: A methodological approach.” *American Journal of Industrial Medicine*, 9(4): 305–321.

Stellman, Steven D., Jeanne Mager Stellman, and John F. Sommer. 1988. “Health and reproductive outcomes among American Legionnaires in relation to combat and herbicide exposure in Vietnam.” *Environmental Research*, 47(2): 150–174.

U.S. Geological Survey (USGS). 1996. “Global 30-Arc-Second Elevation Data, GTOPO30.”

Vuong, Dinh Tuan Nguyen. 2024. “The Persistent Health Effects of Defoliating Vietnam.”

Yamashita, Nobuaki, and Trong-Anh Trinh. 2020. “Long-term effects of Agent Orange on health capital in Vietnam.” , (2020-20).

Yamashita, Nobuaki, and Trong Anh Trinh. 2022. “Long-Term Effects of Vietnam War: Agent Orange and the Health of Vietnamese People After 30 Years.” *Asian Economic Journal*, 36(2): 180–202.

Zabel, Florian, Birgitta Putzenlechner, and Wolfram Mauser. 2014. “Global agricultural land resources: A high resolution suitability evaluation and its perspectives until 2100 under climate change conditions.” *PLoS ONE*, 9(9): 1–12.

ONLINE APPENDIX FOR “NOT GONE WITH THE WIND: LONG-RUN IMPACTS OF HERBICIDAL WARFARE IN VIETNAM”

Gaku Ito* Duc Tran† Yuichiro Yoshida‡

January 26, 2026

Contents

A Data Details and Additional Estimates	A1
A.1 Descriptive Statistics	A1
A.2 Stellman-National Academy of Sciences HERBS File	A1
A.3 Running Variable Coding	A2
A.4 Historical Maps and Archival Sources	A4
A.5 Naive OLS Estimates	A7
B Identification Assumption	A8
B.1 Running Variable Distribution	A8
B.2 Covariate Distribution	A9
C Robustness Check	A13
C.1 Confounding Nonlinearity and Model Dependence	A13
C.2 Flight Buffer Width and Half-Decay Distance	A15
C.3 Counterinsurgency Strategy and Jackknife Estimates	A16
C.4 Unobserved Confounding Forces and Sensitivity Analysis	A18
References	A21

*Associate Professor, Graduate School of Economics, Osaka Metropolitan University. 3-3-138 Sugimoto, Sumiyoshi, Osaka 558-8585, Japan. Email: gaku@omu.ac.jp.  0000-0002-5362-6918. URL: <https://gaku-ito.github.io>. Corresponding author.

†Assistant Professor, Global Campus Institute and Graduate School of Integrated Sciences for Life, Hiroshima University. 1-4-4 Kagamiyama, Higashi-Hiroshima, Hiroshima 739-8528, Japan. Email: tranduc@hiroshima-u.ac.jp.  0000-0002-3443-6510.

‡Professor, School of Economics, Kwansei Gakuin University. 1-155 Uegahara Ichiban-cho, Nishinomiya 662-8501, Japan. Email: yuichiroyoshida@kwansei.ac.jp.  0000-0002-3061-4594.

A Data Details and Additional Estimates

A.1 Descriptive Statistics

Table A.1 reports the summary statistics of the variables used in the empirical analysis. Nonbinary variables excepting the running variable, Viet Cong control prevalence, and population growth rates are log-transformed.

A.2 Stellman-National Academy of Sciences HERBS File

The key source of our empirical analysis is the Stellman-National Academy of Sciences (NAS) version of the Herbicide Report System (HERBS) file (S-NAS HERBS, Stellman et al., 2003a,b). We first web-scraped the records in the database from the Agent Orange Warehouse website.¹ We then converted the geocoordinate information originally recorded in the military grid reference system (MGRS) format into the longitude-latitude format, and combined the converted geocoordinates with other mission- and flight-level information.

In addition to the geocoordinates of spray start, turn, end points, the database contains information about the mission dates, spray methods, and agents and gallons sprayed in individual missions.² The spray method is categorized into “Fixed-Wing Aircraft,” “Ground Spraying,” “Helicopter,” and “Unspecified” categories. As Stellman et al. (2003b, 323) note, a majority of herbicide missions was carried out by fixed-wing aircraft (C-123 airplanes), with 5,961 (65.2%) out of 9,141 missions recorded as aircraft missions. The remaining 3,180 missions include 2,108 (23.1%) missions by helicopters, 446 (4.9%) ground applications, and 626 (6.8%) missions with the method information remaining unspecified. As explained in the main text, our empirical analysis uses the fixed-wing aircraft missions by the US Air Force, which dispersed approximately 95% of all herbicides (Stellman et al., 2003a, 681–682; see also, Figure 1 in the main text).

¹ Available at: <http://www.workerveteranhealth.org/milherbs/new/>, accessed August 29, 2021.

² For the aircraft navigation system in the 1960s and the development of the tactical air navigation system (TACAN) distance measuring equipment (DME) in the region, see, for example, Rowley (1975).

Table A.1: Descriptive Statistics

	Observations	Mean	SD	Median	IQR
Panel A: Outcome Variables					
Population (2001)	716	1.333	0.995	1.193	1.327
Population (2020)	716	1.484	1.109	1.287	1.324
Overall Population Growth Rate (2001–2020)	716	0.152	0.319	0.095	0.338
Annual Population Growth Rate (2001–2020)	13,604	0.008	0.104	0.003	0.074
Panel B: Treatment and Assignment Variables					
HERB ($D^{\text{Half}} = 500\text{m}$)	716	7.093	1.946	7.488	2.411
Edge Distance (in km)	716	-1.140	2.924	-1.418	4.660
$\mathbb{1}[\text{EdgeDist}_{ij} \geq 0]$	716	0.330	0.470	0.000	1.000
Panel C: Covariates					
Key Target Covariates					
North Vietnam Army Base Distance	716	2.758	0.730	2.844	1.059
U.S. Base Distance	716	2.996	0.820	3.149	1.065
U.S. Troop Distance	716	4.892	1.276	5.220	1.778
Population (1967–1969)	716	3.143	3.314	0.693	6.524
Forest Presence	716	0.522	0.500	1.000	1.000
Rice Cropland	716	0.300	0.459	0.000	1.000
Road Distance	716	1.463	1.529	1.671	2.248
Slash and Burn Cropland	716	0.251	0.434	0.000	1.000
Viet Cong Control Prevalence	716	0.172	0.325	0.000	0.129
Geographic Covariates					
Elevation	716	4.319	1.454	3.850	2.918
Flow Accumulation	716	1.747	2.280	0.693	2.996
Precipitation	716	7.464	0.200	7.446	0.315
Rice Suitability	716	1.619	1.322	1.099	3.258
River Distance	716	-0.199	1.336	-0.068	1.719
Ruggedness	716	1.072	0.786	0.693	0.916
u -Component of Neutral Wind					
v -Component of Neutral Wind					
Historical Covariates					
Bombing Point Distance	716	-0.190	0.822	-0.168	1.215
Border Distance	716	4.330	0.828	4.554	0.798
Number of Neighbor Hamlets	716	2.997	0.879	3.045	1.153
Railway Distance	716	2.992	1.632	3.364	2.069

Notes: SD = standard deviation, IQR = interquartile range.

A.3 Running Variable Coding

The main text explains the coding procedure of the running variable with examples in which single flight path polygons cover individual hamlets. When multiple flight polygons cover single hamlets, we follow the same uniform rule for both treated and control hamlets to measure the running variable, as illustrated by Figure A.1. For the treatment group hamlets,

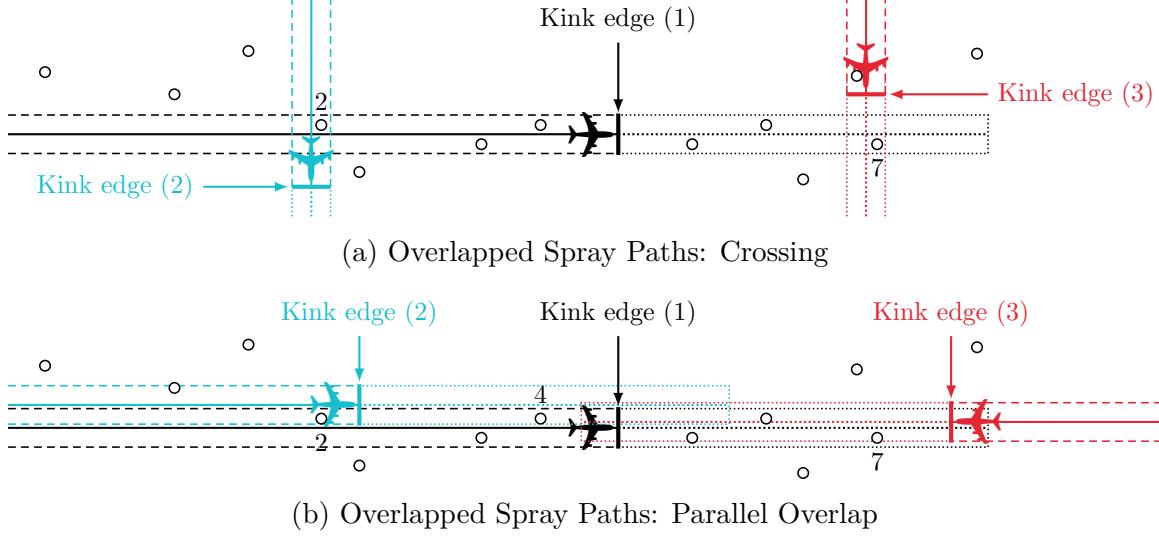


Figure A.1: Running Variable Coding with Overlapped Spray Paths

Notes: (a), (b) Dashed and dotted segments represent distinct spray path polygons. Hamlets 2 and 7 are geographically covered by multiple spray path polygons.

we first take the minimum distance to the kink edges and then assign the maximum, instead of the minimum, of the minimum distances as its running variable. For example, we assign the distance to kink edge (1) instead of the distance to kink edge (2) as the running variable for hamlet 2 (in the treatment group). For the control group hamlets, we first measure the distances to the kink edges of all flight polygons covering the hamlets, and then assign the minimum distance as its running variable. For example, the running variable for hamlet 7 (in the control group) in Figure A.1(a) is measured as the distance to kink edge (3) instead of the distance to kink edge (1).

Note that we follow an identical coding rule to measure the running variable for both treated and control group hamlets where the running variable is defined as the “signed” distance from the cutoff as in literature. Specifically, if a hamlet is always located on the spray-off side (i.e., control group), all distances are negative and the maximum signed distance corresponds to the minimum “absolute” distance to the cutoff (e.g., $\max\{-100, -200, -300\} = -100$). In contrast, if the hamlet is located on the spray-on side for at least one mission (i.e., treated group), the running variable equals the maximum distance among those missions (e.g., $\max\{100, 200, -300\} = 200$). The difference in the coding rule for the treated and

control group hamlets reflects the signs of the running variable that takes a positive value for the treated hamlets and negative values for the control group hamlets.

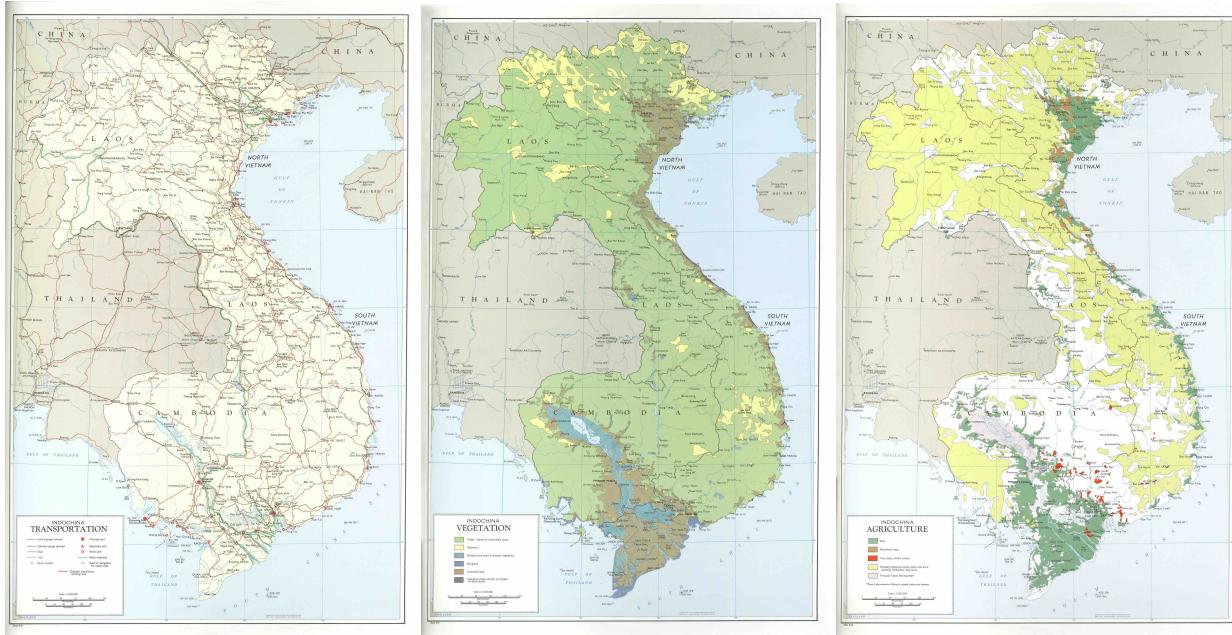
For example, as illustrated by hamlet 2 in Figure A.1(a), when covered by multiple flight path polygons, a treated hamlet may be located in the middle of one flight polygon (e.g., 4 km from kink edge 1) while located closed to the kink edges of another flight path polygon (e.g., 200 meters from kink edge 2). In this case, a simple minimum distance coding assigns 200 meters to hamlet 2 as the running variable. However, this coding significantly underestimates the hamlet's distance to kink edge because the hamlet is located in the middle of another flight path and is strongly exposed to the treatment due to the mission flight with kink edge (1) regardless of its distance to kink edge (2). Our coding rule prevents this underestimation of the running variable and assigns 4 km instead of 200 meters to the treated hamlet. Figure A.1(b) illustrates another example with multiple spray paths with parallel overlaps. The kink edge distances for hamlets 2 and 7 are measured in the same manner as illustrated with the case of Figure A.1(a) (i.e., distance to kink edge 1 for hamlet 2, and distance to kink edge 3 for hamlet 7).

A.4 Historical Maps and Archival Sources

As explained in the main text, some of our covariates relies on historical maps and archival sources. As one of the key sources, we originally georeferenced and image-processed the historical maps of *Indochina Atlas* published by the Central Intelligence Agency (CIA) in 1970 and later digitized by the University of Texas Libraries.³ As shown in Figures A.2(a) to A.2(c), the *Atlas* provides a series of maps as of the period including the transportation lines (i.e., roads, trails, and railways), surface configurations, and crop fields. Figures A.2(d) to (f) show the image-processed results for transportation lines, forest coverage, and croplands.⁴

³The digitized (scanned) maps are available at https://maps.lib.utexas.edu/maps/indochina_atlas/ and <https://www.history.navy.mil/research/library/exhibits/maps/indochina-atlas-1970.html>.

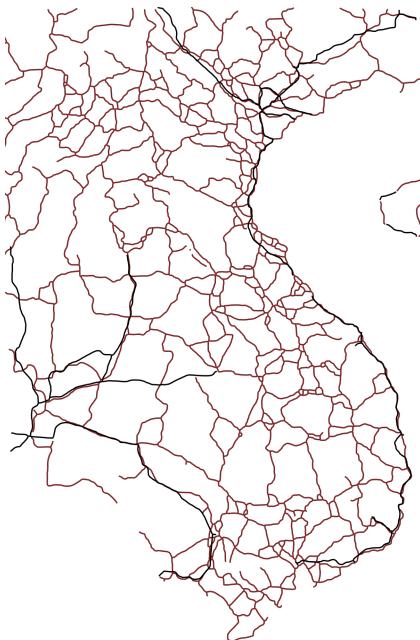
⁴For image-processing, we primarily rely on the mean-shift segmentation algorithm implemented in Orfeo Toolbox (Grizonnet et al., 2017) and the skeletonization algorithm of Tveite (2015). We rely on the same image-processing procedures to generate the polygons of demarcation line and demilitarized zone (DMZ)



(a) Transportation Map

(b) Vegetation Map

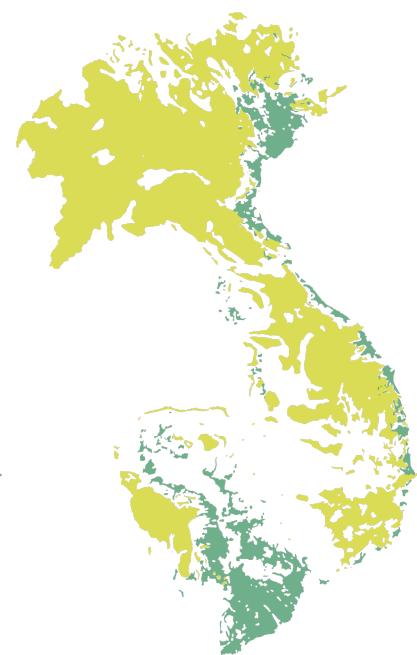
(c) Agriculture Map



(d) Road and Railway Network



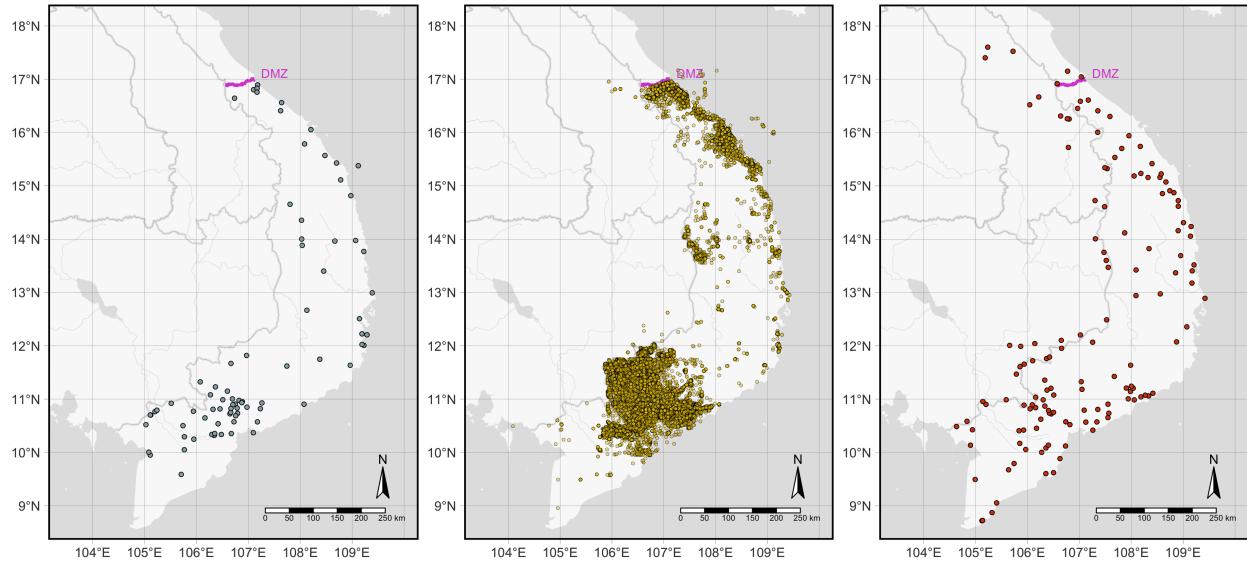
(e) Forest Coverage



(f) Rice Croplands and Slash and Burn Cultivation Areas

Figure A.2: *Indochina Atlas* Maps and the Image-Processed Vector Objects

Notes: Panels (a), (b), and (c) depict the original (non-georeferenced) scanned maps of *Indochina Atlas*, digitized by the University of Texas Libraries. Panels (d), (e), and (f) show the generated vector objects (lines and polygons). Red and black segments in Panel (d), respectively, indicates roads (including trails) and railways. Green region in Panel A.2(e) represents forest presence. In Panel (f), green regions represent rice croplands, and yellow regions indicate the areas with slash and burn cultivation.



(a) U.S. Air Force and Navy Bases (b) U.S. Army and Marine Troops (c) Suspected North Vietnam Army (NVA) Base Areas

Figure A.3: Locations of U.S. Bases, Troops, and NVA Base Areas

Notes: Dots in each panel indicate, respectively, (a) U.S. Air Force and Navy bases, (b) U.S. Army and Marine troop locations, and (c) suspected areas of North Vietnam Army (NVA) base areas in and near South Vietnam. Dots in (c) represent the geocoordinates of the “approximate center” of the suspected NVA base areas (National Archives and Records Administration. 2007. “Reference Copy of Technical Documentation for Accessioned Electronic Records: Enemy Base Area File (BASFA), 7/1967–7/1971, Translation File, RG 330, Records of the Office of the Secretary of Defense.” pp.26–27. Available at: <https://catalog.archives.gov/id/2573252>, accessed June 16, 2023).

We rely on several other sources to measure the historical landscape of military bases and troops. Suspected areas of North Vietnam Army (NVA) base are retrieved from the Enemy Base Area File (BASFA), July 1, 1967–July 1, 1971 hosted by the National Archives and Records Administration (NARA).⁵ The locations of U.S. Air Force and Navy bases and U.S. Army and Marine troops are retrieved from the S-NAS HERBS database. Figure A.3 depicts the spatial distributions of U.S. bases, U.S. troop locations, and the suspected areas of NVA bases in and near South Vietnam.

from CIA’s “Vietnam demarcation line and demilitarized zone. 12-66” map as of 1966 (<https://www.loc.gov/resource/g8021f.ct002840/>, accessed September 22, 2021).

⁵National Archival Identifier 2573252; Accession Number of NN3-330-76-037. Available at: <https://catalog.archives.gov/id/2573252>, accessed June 16, 2023.

A.5 Naive OLS Estimates

Table A.2 reports the uninstrumented, naive ordinary least square (OLS) version of the main fuzzy regression kink (RK) specification with population size in 2001 and 2020 as outcomes. As briefly reported in the main text, compared with the main RK estimates, the naive OLS estimates underestimate or even fail to detect the negative association between herbicide exposure and contemporary population size revealed in Table 1.

The RK-OLS discrepancy may reflect (1) bias in the OLS estimates induced by confounding bias and measurement error, (2) bias in the RK estimates, (3) bias in both estimates, and (4) the difference in the estimands (average treatment effect, ATE, and the LAR or TT effect). Given the nonrandom nature of herbicide assignment, we mainly attribute the difference between the OLS and RK estimates to the bias remaining in the OLS estimates.

Table A.2: Herbicide Exposure and Contemporary Population Size, Naive OLS Estimates

Panel A: ln Population (2001)					
	(1)	(2)	(3)	(4)	(5)
ln HERB	-0.074** (0.036)	-0.039 (0.031)	-0.049 (0.031)	-0.052* (0.030)	-0.051* (0.030)
Adjusted R ²	0.627	0.735	0.758	0.765	0.769
Panel B: ln Population (2020)					
	(1)	(2)	(3)	(4)	(5)
ln HERB	-0.089** (0.040)	-0.057 (0.035)	-0.066* (0.034)	-0.069** (0.033)	-0.068** (0.033)
Adjusted R ²	0.639	0.730	0.753	0.757	0.760
Observations	519	519	519	519	519
Avg. N neighbors (Conley SE cluster)	27.5	27.5	27.5	27.5	27.5
Key Target Covariates		✓	✓	✓	✓
Geographic Covariates			✓	✓	✓
Historical Covariates				✓	✓
ln Spatially-Lagged HERB					✓
Fixed effects and $f(\text{Lon}, \text{Lat})$	✓	✓	✓	✓	✓

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Conley (1999) standard errors adjusted for spatial clustering with a 30 km cutoff and a Bartlett kernel are in parentheses. **Key target covariates**: NVA base distance, population (1967–1969), U.S. base distance, U.S. troop distance, forest presence, rice cropland, road distance, slash and burn cropland, Viet Cong control prevalence. **Geographic covariates**: Precipitation, temperature, u - and v -components of neutral wind, elevation, flow accumulation, rice suitability, river distance, ruggedness. **Historical covariates**: Bombing point distance, border distance, number of neighbor hamlets, railway distance. ln spatially-lagged HERB is the logged average HERB among the neighbor hamlets with a 30 km cutoff. **Fixed effects**: Agent fixed effect, province fixed effect, end-leg fixed effect, mission phase fixed effect.

B Identification Assumption

B.1 Running Variable Distribution

A valid fuzzy RK design hinges on the smoothness assumption, which yields several testable implications (Card et al., 2015b). The first testable implication of the smoothness assumption requires a running variable continuously differentiable at the kink point. To empirically validate the implication, Figure B.1 examines the continuity of the running variable using the estimator of Cattaneo et al. (2020). Consistent with the identification assumption, the density test fails to detect discernible or statistically significant discontinuity in the running variable at the kink point ($t = -0.209$ and $p = 0.835$).

Second, the smoothness assumption requires not only the absence of discontinuity but also the absence of a slope change in the running variable distribution at the kink point. To further validate the smoothness assumption, we follow the approach of Bana et al. (2020), Card et al. (2015a,b), and Landais (2015) and test for the kink in the running variable distribution via polynomial regressions. Formally, and with abuse of notation, we first aggregate the hamlet observations using bins with different sizes based on the running variable and then estimate the following regression model:

$$N_b^{\text{obs.}} = \beta \mathbb{1}[\overline{\text{EdgeDist}}_b \geq 0] + \sum_{p=1}^P \left[\gamma_p \overline{\text{EdgeDist}}_b^p + \delta_p \overline{\text{EdgeDist}}_b^p \cdot \mathbb{1}[\overline{\text{EdgeDist}}_b \geq 0] \right] + e_b, \quad (\text{B.1})$$

where b indexes bins, $N^{\text{obs.}}$ reflects the number of observations in each bin, $\overline{\text{EdgeDist}}$ is the midpoint of EdgeDist of each bin, and P is the polynomial order. The coefficient on the interaction between the linear of the running variable and the treatment group indicator, δ_1 , captures the change in the slope of the probability density function of the running variable at the kink point. We repeatedly estimate the polynomial regression using different bin widths ranging from 50 meters to 500 meters with an increment of 50 meters and the running variable polynomial orders ranging from 2 to 6.

Figure B.2 summarizes the polynomial regression results by plotting the coefficient es-

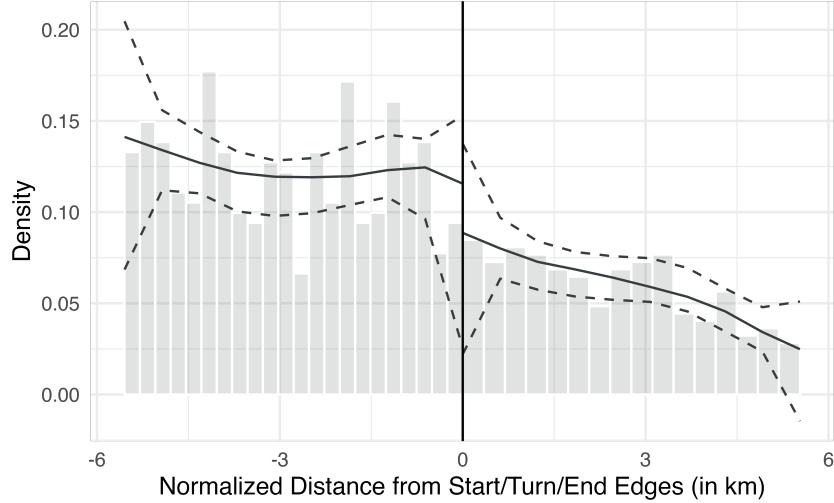


Figure B.1: Continuity of the Running Variable Distribution

Notes: Gray bars represent the histogram estimates of the running variable, EdgeDist. Solid and dashed lines indicate the local polynomial density estimates and the corresponding 95% confidence intervals (Discontinuity test: $t = -0.209$ and $p = 0.835$). The discontinuity estimate is computed using the procedure of Cattaneo et al. (2020), with the running variable recentered at zero.

timates on the interaction term between the treatment group indicator and the (linear) running variable (δ_1). Triangles in each panel and Panel (f) indicate the estimates obtained from the smallest Akaike Information Criterion (AIC) value for each bin size. Regardless of the bin size and polynomial order, the coefficient estimates remain statistically insignificant at the 5% level. The polynomial regressions reveal no discernible slope change in the running variable distribution at the kink point and fail to invalidate the current RK strategy.

B.2 Covariate Distribution

Another key testable implication of the smoothness assumption involves covariate distributions around the kink point. The validity of the current RK design would be undermined if predetermined covariates, along the treatment, exhibit discernible kink around the cutoff.⁶ We examine the validity of the assumption by subsequently estimating the first-stage spec-

⁶Seen from the instrumental variable (IV) perspective, a key assumption here is that the instrument, $\text{EdgeDist} \times 1[\text{EdgeDist} \geq 0]$, affects the outcome (slope change) only through the treatment (slope change). A kink in a covariate can also be viewed as a sign of exclusion restriction violation such that the instrument affects the outcome through the covariate, not exclusively through the treatment.

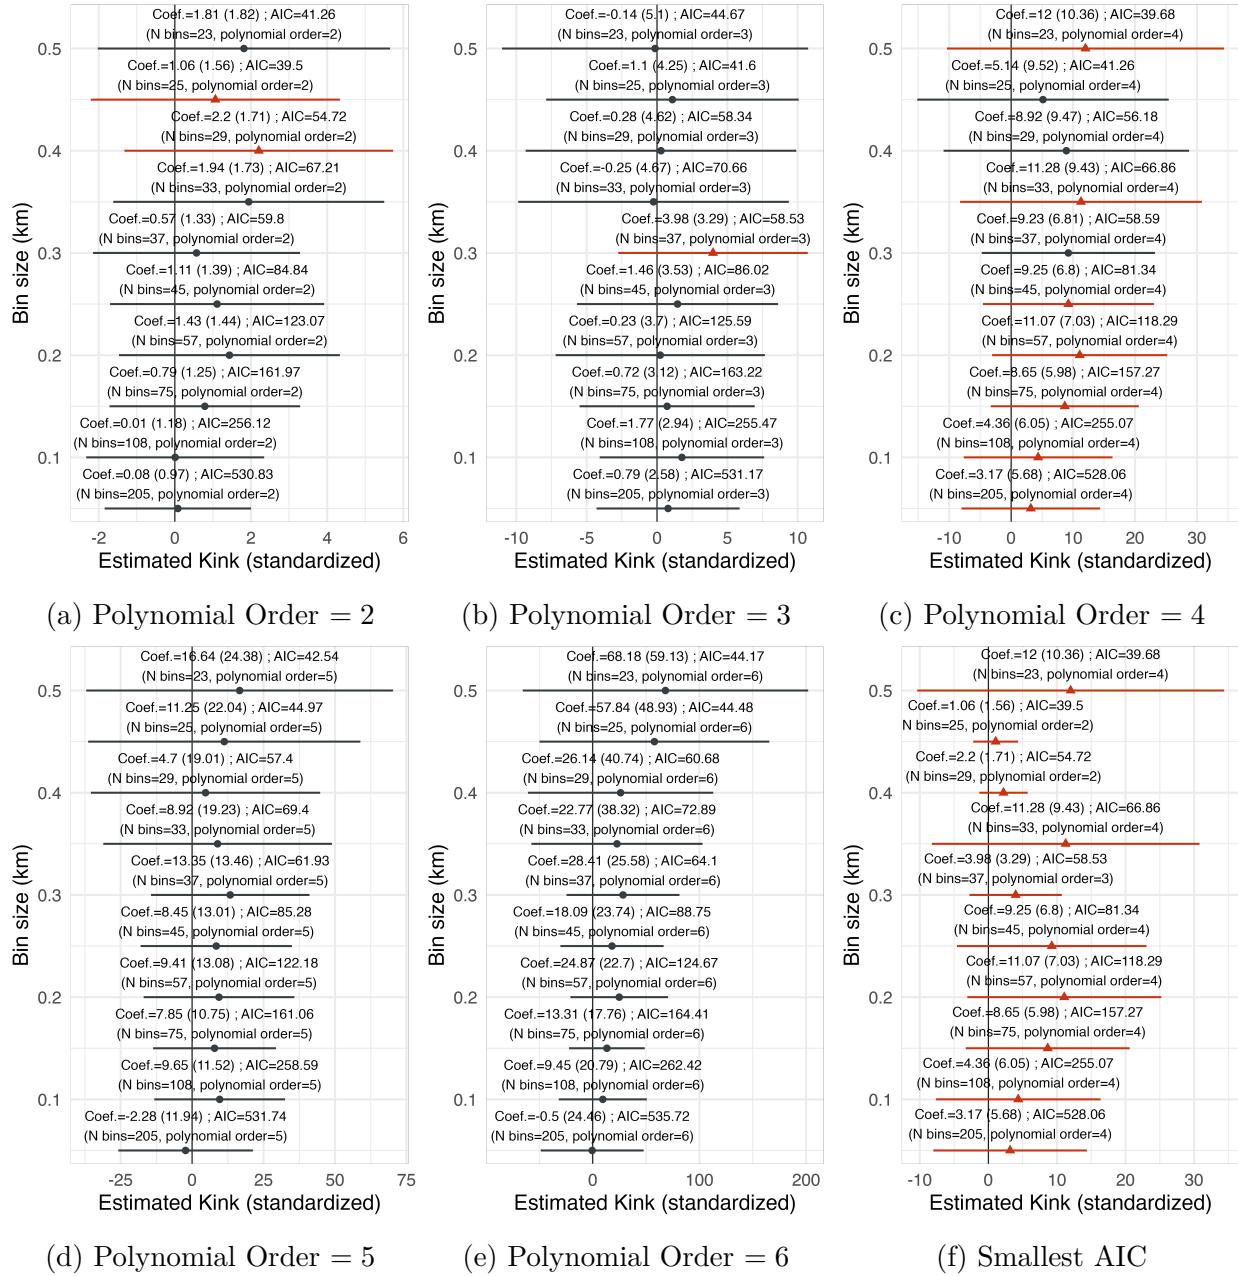


Figure B.2: Kink Estimates in the Running Variable Distribution

Notes: Dots represent the kink estimates obtained from polynomial regressions (δ_1 in equation B.1), with the running variable polynomial order specified as in the figure labels and bin sizes on the vertical axis. Triangles represent the estimates obtained from the model with the smallest Akaike Information Criterion (AIC) for each bin size. Horizontal segments represent the 95% confidence intervals.

ification of equation (3) with the left-hand-side variable, $\ln \text{HERB}$, replaced by one of the covariates. We estimate the placebo kink estimate with and without adjustments for the remaining covariates, with the baseline bandwidth $b = 4$ km as in Table 1 in the main text. The

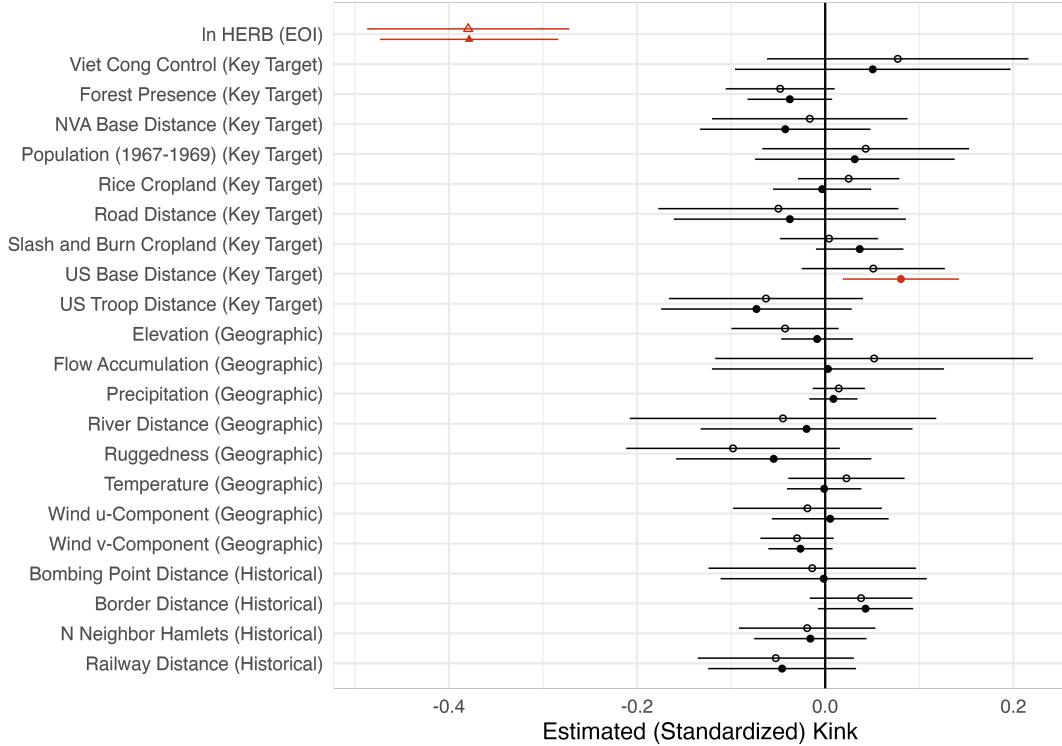


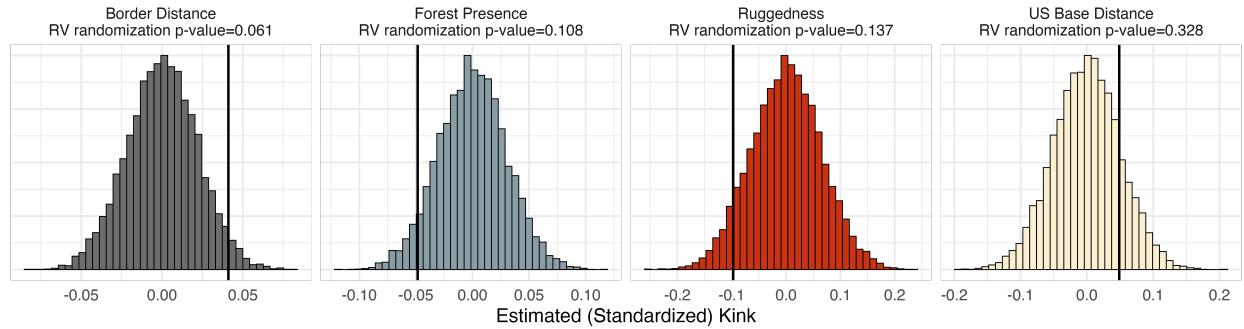
Figure B.3: (Standardized) Covariate and Treatment Kink Estimates

Notes: Symbols represent the reduced-form kink estimates with the variable on the vertical axis as the outcome variables, with (solid) and without adjustments for the remaining covariates (hollow). Horizontal segments represent 95% confidence intervals based on [Conley \(1999\)](#) standard errors with a 30 km cutoff and a Bartlett kernel (not corrected for multiple testing). All nonbinary variables are standardized.

smoothness assumption requires that the coefficient on $\text{EdgeDist} \times \mathbb{1}[\text{EdgeDist} \geq 0]$ remains indistinguishable from zero once we replace the left-hand-side variable with covariates.

Figure B.3 summarizes the (standardized) kink estimates with (hollow symbols) and without adjustments for the remaining covariates (solid symbols), along with the treatment kink estimates (top row triangles). In a sharp contrast with the treatment slope change, and consistent with the identification assumption, most of the placebo regressions fails to detect a discernible slope change around the kink point. Exceptions include the substantively small but statistically significant estimates for U.S. base distance (coef. = 0.078, $t = 2.559$) and the marginally significant estimates for border distance (coef. = 0.043, $t = 1.657$) when adjusting for the remaining covariates (solid symbols), terrain ruggedness (coef. = -0.098 , $t = -1.696$) without covariate adjustments (hollow symbols), and forest presence

RV Randomization, without Covariate Adjustment



RV Randomization, with Covariate Adjustment

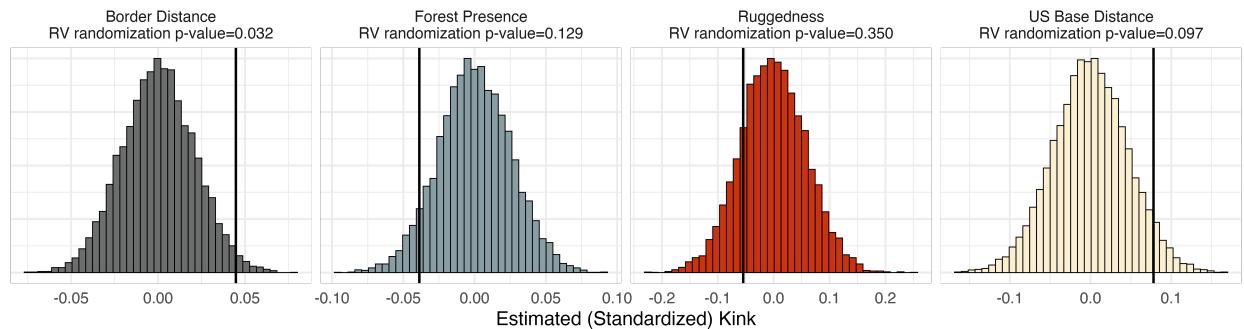


Figure B.4: Running Variable Randomization Inference

Notes: Each panel shows the histogram of 10,000 placebo kink estimates for the outcome in the top label. We randomly assign (permute) the running variable to the sample hamlets, and then estimate the reduced-form (or first-stage) version of the RK specification with (bottom row panels) and without covariate adjustments (top row panels) for 10,000 times. Solid vertical segments indicate the actual kink estimates with the observed running variable. The randomization (two-sided) p -values in the top labels indicate the share of the absolute placebo coefficients that are larger than the absolute actual coefficient.

(coef. = $-0.048, t = 1.623$ without covariate adjustments and coef. = $-0.038, t = -1.638$ with covariate adjustments).⁷

Nonetheless, the running variable (RV) randomization exercise disagrees with asymptotic inference and fails to negate that the covariate kink are generated by chance. As in Table 1, we randomly assign the running variable to the sample hamlets and estimate the placebo kink coefficient for 10,000 times. The randomization two-sided p -value indicates the share of the absolute placebo coefficients that are larger than the absolute actual coefficient. Figure B.4 shows the histograms of the (standardized) placebo estimates and the corresponding actual estimates (vertical lines), with the corresponding randomization p -values larger than 0.05

⁷Given one treatment and 22 covariates (23 variables) in Figure B.3, the probability of a Type I error in which one falsely rejects at least one null hypothesis with $\alpha = 0.05$ is 69.3% ($1 - 0.95^{23} = 0.693$).

excepting for border distance with covariate adjustments (randomization $p = 0.032$).

A natural and substantively important question is on how the covariate kink affects the main findings. A focused exercise in Appendix C, however, suggests that the RK estimates remain stable across all possible model specifications (covariate combinations) and outcomes, and reveals little sign of model dependence. Along with the little evidence of model dependence, the covariate kink test fails to invalidate the current RK design.

C Robustness Check

C.1 Confounding Nonlinearity and Model Dependence

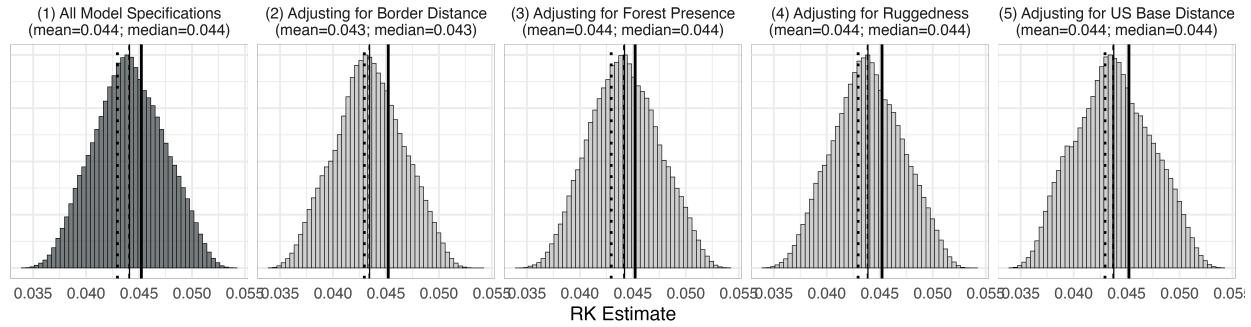
One may wonder how the kink in covariates and model specification influence the main findings. Related to this point, [Ando \(2017\)](#) underlines the susceptibility of RK estimates to the failure to adjust for confounding nonlinearity. A simple, crude way to address this concern for confounding nonlinearity is to estimate *all* possible model specifications to generate the empirical distributions of the coefficient estimates, and compare the reported estimates with the resultant distributions. The underlying idea is that if the current RK design suffers confounding kink, the estimates should be sensitive to the choice of covariate adjustments.

Specifically, we repeatedly estimate the reduced-form model of Equation (2) in the main text with each covariate combination and three cross-sectional outcomes.⁸ Given that our covariate set includes 22 covariates along with the fixed effects and the longitude-latitude polynomial, the number of possible model specifications amounts to $2^{21} = 2,097,152$ for each outcome, treating the u - and v -components of neutral winds as a single block and omitting the spatially-lagged EOI and interaction terms.⁹ We estimate each model specification using the sample hamlet with the baseline bandwidth of $b = 4$ km as in Table 1 in the main text.

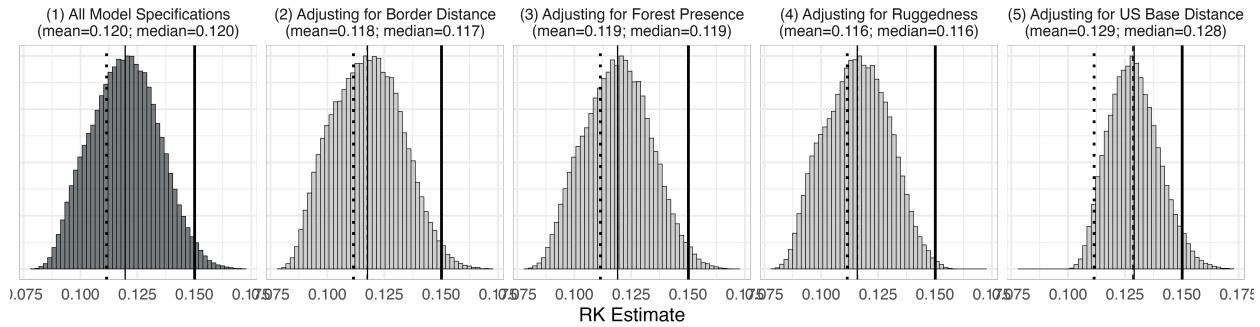
⁸This exercise relies on the cross-sectional setup given the qualitative similarity of annual (panel) and overall (cross-sectional) population growth rate estimates in Figure 4 in the main text.

⁹To improve computational efficiency, we always jointly include or exclude the u - and v -components of neutral winds, while treating all remaining covariates independently.

Outcome: Overall Population Growth Rate (2001-2020)



Outcome: In Population (2001)



Outcome: In Population (2020)

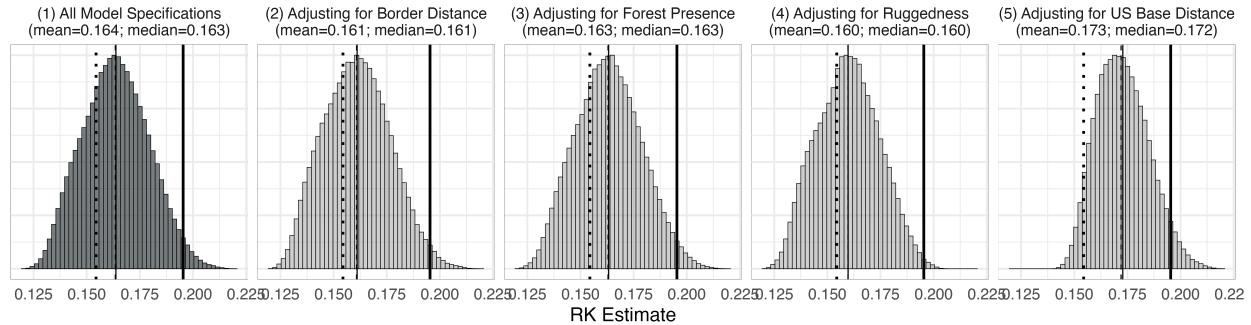


Figure C.1: Empirical Distribution of Second-Stage RK Estimates, All Possible Specifications

Notes: Density histograms in each row plot the empirical distributions of (subsets of) $2^{21} = 2,097,152$ RK model specifications for three cross-section outcomes, with different covariate adjustment conditions denoted in the top text labels. Thin solid and dashed vertical segments indicate the mean and median estimates in each panel. Bold solid and dotted vertical segments represent the reduced-form estimates with the specifications of Model 1 (without covariates, solid) and Model 4 (with all covariates, dotted) in Table 1 in the main text.

Figure C.1 summarizes the result of the “try-all” exercise. To explicitly examine how the adjustments for the variables in Figure B.4 affect the estimates, Figure C.1 presents density histograms for RK estimates of, from the left, (1) all possible 2,097,152 models, (2) $2^{20} = 1,048,576$ models adjusting for border distance, (3) 1,048,576 models adjusting for forest

presence, (4) 1,048,576 models adjusting for terrain ruggedness, and (5) 1,048,576 models adjusting for U.S. base distance for each outcome. Thin solid and dashed vertical segments indicate the corresponding mean and median estimates of the empirical distributions, and bold solid and dotted segments indicate the reduced-form estimates of Model 1 (without covariates) and Model 4 (adjusting for all covariates) in Table 1 in the main text.

The key results are threefold. First, across the outcomes, the empirical distribution of RK estimates is close to a normal distribution (with slight skews in population size in 2001 and 2020), suggesting that covariate combination does not play a major role in altering the RK estimates beyond random error. Second, the adjustment for border distance and other covariates is unlikely to alter the RK estimate. As graphically illustrated by the similarity between the panels in each row, the RK estimate distribution remains almost unchanged regardless of covariate adjustments, with almost identical sample mean (solid vertical segment) and median estimates (dashed segment). Third, the baseline reduced-form estimates (indicated by the bold solid and dotted segments) align closely with the means and medians of the empirical coefficient distributions. Specifically, models with the full set of covariates (dotted segments) slightly underestimate the reduced-form kink, whereas models without covariate adjustments (bold solid segments) tend toward overestimation. The mean and median of the distribution are bracketed by these two specifications.

Combined, these results suggest that the covariate kink and the choice of covariate adjustments are unlikely to affect the inference severely, and the main findings are unlikely to be a product of arbitrary model picking or (observed) confounding nonlinearity.

C.2 Flight Buffer Width and Half-Decay Distance

Admittedly, the main RK estimates rely on an arbitrary combination of two parameters, flight buffer width and half-decay distance, to obtain the RK sample and measure the hamlet-level herbicide exposure score, HERB. As explained in Section 4 of the main text and Section A.3, flight buffer width defines the sample hamlets used in the RK estimation,

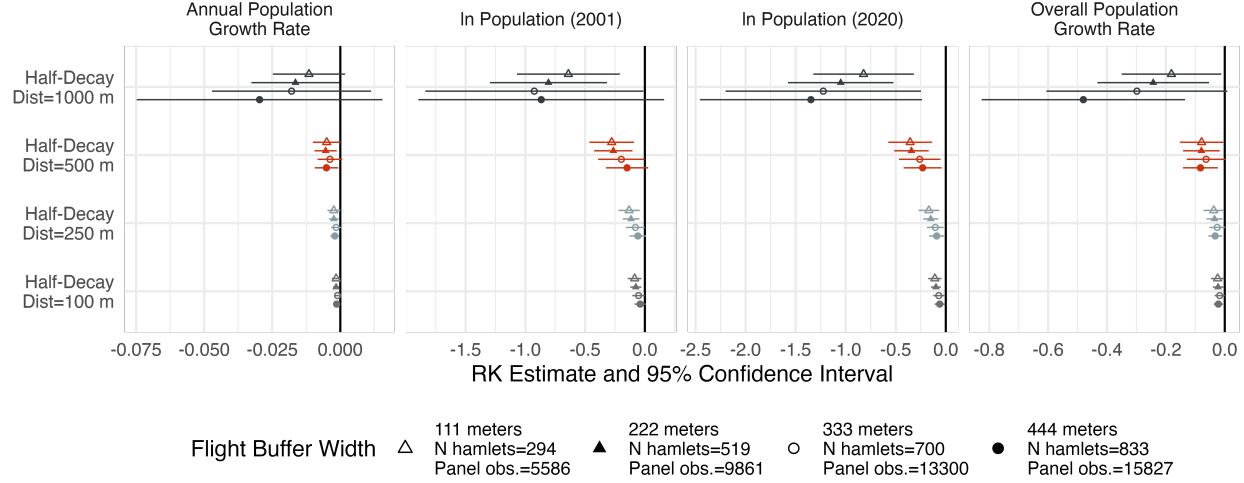


Figure C.2: RK Estimates with Different Flight Buffer Width and Half-Decay Distance Settings

Notes: Symbols represent the second-stage RK estimates with different flight buffer widths (indicated by symbols) and the half-decay distance parameter on the vertical axis, for each of the outcome variables in the column labels. Horizontal segments display the 95% confidence intervals based on Conley (1999) standard errors with a 30 km cutoff and a Bartlet kernel for the cross-sectional specification and hamlet-level clustering for the panel setup. Red solid triangles with half-decay distance = 500 m indicate the baseline estimates reported in Table 1 and Figure 4 in the main text. All estimates rely on the baseline bandwidth of 4 km and the model specification of model (1) in Table 1. Panel specification further adjusts for a year fixed effect.

and the half-decay distance parameter determines HERB (Equation 1 in the main text).

To examine the robustness to the particular parameter setting, Figure C.2 reestimates the RK model with alternative parameter combinations. While the coefficient size varies, the negative association between herbicide exposure and contemporary population outcomes remains robust across different RK samples and distance-decay parameter values.

C.3 Counterinsurgency Strategy and Jackknife Estimates

A remaining but important concern for the current RK design is that the results might be driven by *unobserved* confounding kink such as counterinsurgency strategies and authorization procedure of herbicide missions. During the Vietnam War, the U.S. military divided South Vietnam into four Corps Tactical Zones (CTZs) commanded by different military branches and implemented distinct counterinsurgency strategies in individual CTZs. Dell and Querubin (2018, 51–54) exploit the CTZ I–II boundary as one of the sources for causal

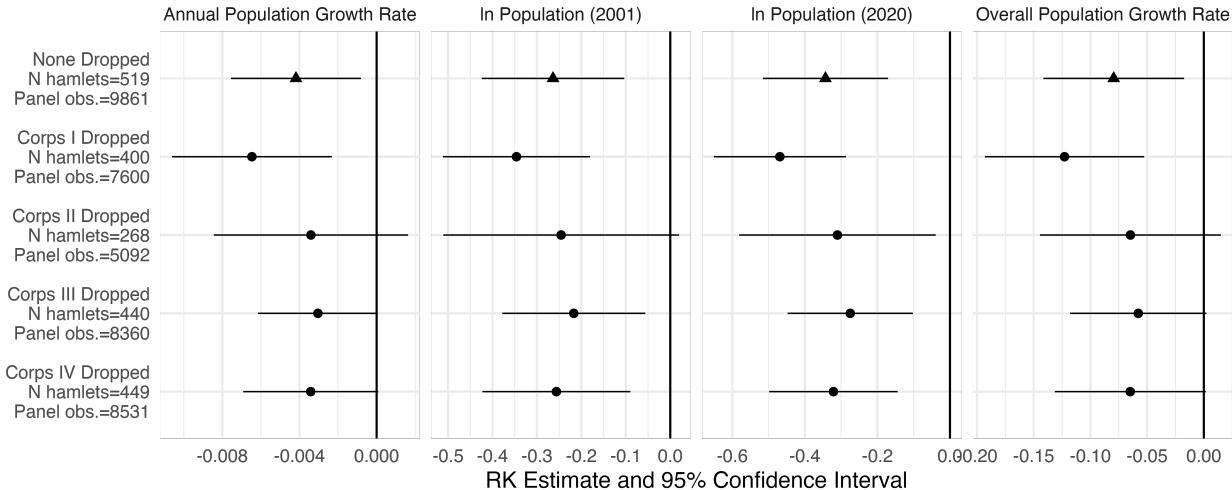


Figure C.3: Corps Tactical Zone-Level Jackknife

Notes: Symbols represent the RK estimates with subsamples dropping the hamlets in the Corps Tactical Zone on the vertical axis. Horizontal segments display the corresponding 95% confidence intervals based on standard errors with a 30 km cutoff and a Bartlet kernel for the cross-sectional specification and hamlet-level clustering for the panel setup. All estimates rely on the baseline bandwidth of 4 km and the specification of model (1) in Table 1. Panel specification further adjusts for a year fixed effect.

identification, given the historical fact that the U.S. Marine Corps (USMC) in Corps Region I implemented a moderate hearts-and-minds-oriented strategy and small-unit operations, while the U.S. Army commanded Corps Region II and emphasized overwhelming firepower and large-scale operations. Similarly, as briefly explained in the main text, the authorization process of herbicide missions involved the South Vietnam Government, both at the national and local levels, and the U.S. at the levels of the U.S. Ambassador, the commander of the U.S. Military Assistance Command, Vietnam (MACV), and the CTZ commanders (Buckingham, 1982, 37; IoM, 1995, 86).

Although the baseline specification includes a district fixed effect, we rely on CTZ-level and province-level jackknife approaches to more explicitly gauge the impact of the CTZs and corresponding differences in counterinsurgency strategies have on the reported RK estimates. Figure C.3 reports the jackknife estimates using a subsample dropping the hamlet observations located in the CTZ on the vertical axis. As the CTZ boundaries follow the preexisting provincial boundaries, and given the role of province chiefs in Ranch Hand mis-

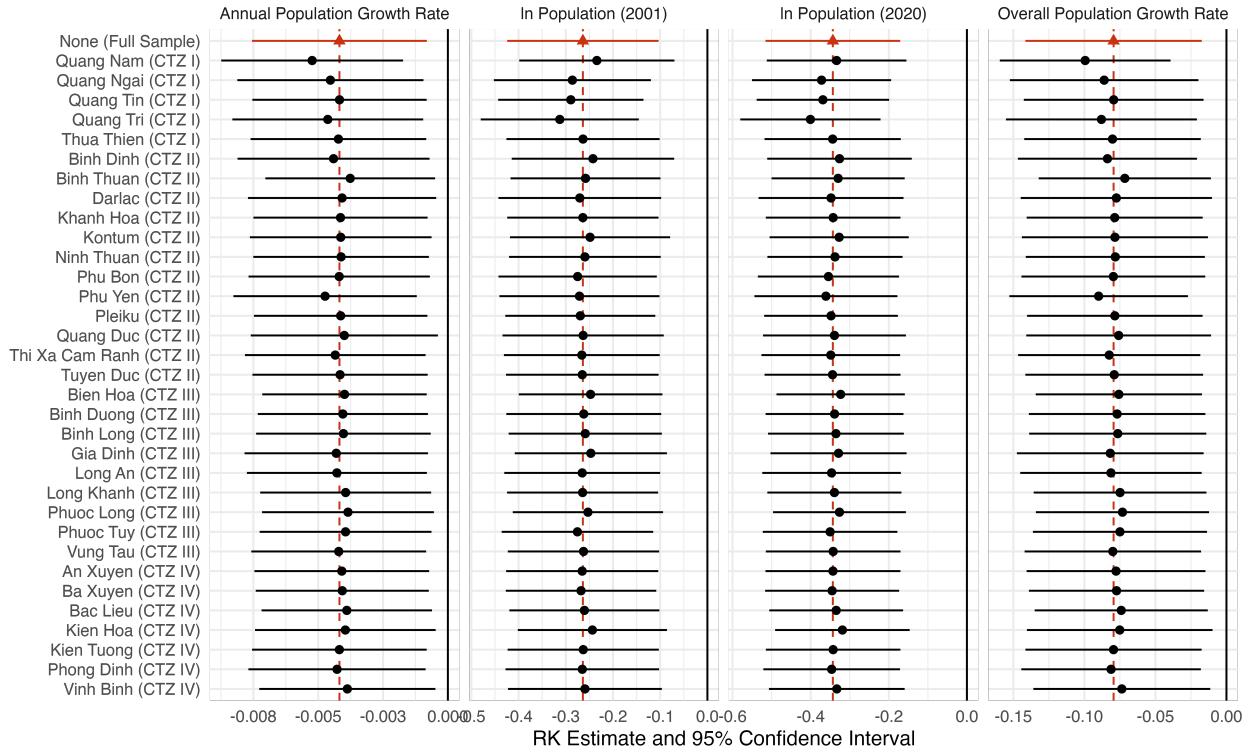


Figure C.4: Province-Level Jackknife

Notes: Symbols represent the RK estimates with subsamples dropping the hamlets in the province on the vertical axis. Horizontal segments display the corresponding 95% confidence intervals based on standard errors with a 30 km cutoff and a Bartlet kernel for the cross-sectional specification and hamlet-level clustering for the panel setup. Triangles and vertical dashed segments indicate the baseline estimate in Table 1 and Figure 4 in the main text. All estimates rely on the baseline bandwidth of 4 km and the model specification of model (1) in Table 1. Panel specification further adjusts for a year fixed effect.

sions, Figure C.4 presents a similar jackknife exercise in which hamlet observations in the province on the vertical axis dropped from the subsample for RK estimate. While the sample size and statistical significance vary, the negative coefficient estimates remain qualitatively unchanged across subsamples.

C.4 Unobserved Confounding Forces and Sensitivity Analysis

We also use a sensitivity analysis approach to examine how sensitive the main findings are to unobserved confounding nonlinearity. Table C.1 reports the reduced-form and first-stage versions of the main RK specification in Table 1 in the main text. The “robustness values” in each panel indicate the percentage of residual variance that unobserved confounders would

Table C.1: Reduced-Form and First-Stage Estimates with Robustness Values

Panel A: ln Population (2001)					
	(1)	(2)	(3)	(4)	(5)
EdgeDist $\times \mathbb{1}[\text{EdgeDist} \geq 0]$	0.150*** (0.046)	0.145*** (0.042)	0.119*** (0.039)	0.111*** (0.040)	0.107*** (0.040)
Robustness Value	43.1%	45.2%	41.6%	38.49%	37.4%
Adjusted R ²	0.627	0.740	0.760	0.766	0.770
Panel B: ln Population (2020)					
	(1)	(2)	(3)	(4)	(5)
EdgeDist $\times \mathbb{1}[\text{EdgeDist} \geq 0]$	0.195*** (0.050)	0.191*** (0.045)	0.160*** (0.042)	0.154*** (0.044)	0.150*** (0.044)
Robustness Value	49.25%	51.6%	48.21%	45.71%	44.79%
Adjusted R ²	0.641	0.736	0.756	0.759	0.762
Panel C: ln HERB					
	(1)	(2)	(3)	(4)	(5)
EdgeDist $\times \mathbb{1}[\text{EdgeDist} \geq 0]$	-0.569*** (0.081)	-0.552*** (0.075)	-0.565*** (0.072)	-0.564*** (0.070)	-0.564*** (0.070)
Robustness Value	69.02%	70.53%	72.69%	73.31%	73.49%
Adjusted R ²	0.570	0.609	0.613	0.617	0.617
Observations	519	519	519	519	519
Avg. N neighbors (Conley SE cluster size)	27.5	27.5	27.5	27.5	27.5
Key Target covariates		✓	✓	✓	✓
Geographic covariates			✓	✓	✓
Historical covariates				✓	✓
ln spatially-lagged HERB					✓
Fixed effects and $f(\text{Lon}, \text{Lat})$	✓	✓	✓	✓	✓

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Conley (1999) standard errors adjusted for spatial clustering with a 30 km cutoff and a Bartlett kernel are in parentheses. **Key target covariates**: NVA base distance, population (1967–1969), U.S. base distance, U.S. troop distance, forest presence, rice cropland, road distance, slash and burn cropland, Viet Cong control prevalence. **Geographic covariates**: Precipitation, temperature, u - and v -components of neutral wind, elevation, flow accumulation, rice suitability, river distance, ruggedness. **Historical covariates**: Bombing point distance, border distance, number of neighbor hamlets, railway distance. ln spatially-lagged HERB is the logged average HERB among the neighbor hamlets with a 30 km cutoff. **Fixed effects**: Agent fixed effect, province fixed effect, end-leg fixed effect, mission phase fixed effect. Robustness values indicate the percentage of residual variance that unobserved confounders would need to explain in order to eliminate the reported kink coefficients.

need to explain away the reported kink estimates (Cinelli and Hazlett, 2020). For both reduced-form and first-stage estimates, the robustness values suggest only implausibly strong unobserved confounders are capable of eliminating the main RK estimates.¹⁰ More precisely,

¹⁰We report reduced-form and first-stage estimates as the robustness value measures the strength of unobserved confounding by partial R^2 . Note also that the RK estimand, τ_{RK} , can be written as the ratio of the reduced-form slope change (i.e., outcome slope change at the kink point) relative to the first-stage slope change (i.e., treatment slope change at the kink point) as (Card et al., 2015b):

$$\tau_{\text{RK}} = \left(\frac{\lim_{v_0 \downarrow 0} \frac{d(\mathbb{E}[Y|V=v])}{dv} \Big|_{v=v_0}}{\lim_{v_0 \uparrow 0} \frac{d(\mathbb{E}[Y|V=v])}{dv} \Big|_{v=v_0}} - \frac{\lim_{v_0 \uparrow 0} \frac{d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0}}{\lim_{v_0 \downarrow 0} \frac{d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0}} \right) / \left(\frac{\lim_{v_0 \downarrow 0} \frac{d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0}}{\lim_{v_0 \uparrow 0} \frac{d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0}} - \frac{\lim_{v_0 \uparrow 0} \frac{d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0}}{\lim_{v_0 \downarrow 0} \frac{d(\mathbb{E}[D|V=v])}{dv} \Big|_{v=v_0}} \right),$$

the robustness values indicate that unobserved confounders would need to explain at least 37.4% (Model 5, Panel A) or 44.79% (Model 5, Panel B) of the residual variance of both the instrument ($\text{EdgeDist} \times \mathbb{1}[\text{EdgeDist} \geq 0]$) and the outcome (ln population in 2001 or 2020) to eliminate the reduced-form association. Similarly, unobserved confounders would need to explain at least 69.02% (Model 1, Panel C) of the residual variance of the instrument and the EOI (ln HERB) to eliminate the first-stage association. Given the battery of covariate adjustments in addition to the geocoordinate polynomial and fixed effects, it is implausible that a (combination of) confounding kink is driving the finding.

where V denotes the running variable, v_0 the kink point, and D the treatment. The reduced-form (first-stage) robustness value can be interpreted as a sensitivity measurement for the numerator (denominator).

References

Ando, Michihito. 2017. "How much should we trust regression-kink-design estimates?" *Empirical Economics*, 53(3): 1287–1322.

Bana, Sarah H., Kelly Bedard, and Maya Rossin-Slater. 2020. "The Impacts of Paid Family Leave Benefits: Regression Kink Evidence from California Administrative Data." *Journal of Policy Analysis and Management*, 39(4): 888–929.

Buckingham, William A. 1982. *Operation Ranch Hand: The Air Force and Herbicides in Southeast Asia, 1961-1971*. Washington, DC: Office of Air Force History.

Card, David, Andrew Johnston, Pauline Leung, Alexandre Mas, and Zhuan Pei. 2015a. "The effect of unemployment benefits on the duration of unemployment insurance receipt: New evidence from a regression kink design in Missouri, 2003-2013." *American Economic Review*, 105(5): 126–130.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber. 2015b. "Inference on Causal Effects in a Generalized Regression Kink Design." *Econometrica*, 83(6): 2453–2483.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma. 2020. "Simple Local Polynomial Density Estimators." *Journal of the American Statistical Association*, 115(531): 1449–1455.

Cinelli, Carlos, and Chad Hazlett. 2020. "Making sense of sensitivity: Extending omitted variable bias." *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 82(1): 39–67.

Conley, Timothy. G. 1999. "GMM estimation with cross sectional dependence." *Journal of Econometrics*, 92(1): 1–45.

Dell, Melissa, and Pablo Querubin. 2018. "Nation Building Through Foreign Intervention: Evidence from Discontinuities in Military Strategies." *Quarterly Journal of Economics*, 133(2): 701–764.

Grizonnet, Manuel, Julien Michel, Victor Poughon, Jordi Inglada, Mickaël Savinaud, and Rémi Cresson. 2017. "Orfeo ToolBox: Open source processing of remote sensing images." *Open Geospatial Data, Software and Standards*, 2(1): 15.

Institute of Medicine (Committee to Review the Health Effects in Vietnam Veterans of Exposure to Herbicides). 1995. *Veterans and Agent Orange: Health Effects of Herbicides Used in Vietnam*. Washington, D.C.: National Academy Press.

Landais, Camille. 2015. "Assessing the welfare effects of unemployment benefits using the regression kink design." *American Economic Journal: Economic Policy*, 7(4): 243–278.

Rowley, Ralph A. (Office of Air Force History). 1975. "The Air Force in Southeast Asia: FAC Operations, 1965–1970." *Office of Air Force History, May 1975, available at: <https://media.defense.gov/2011/Mar/24/2001330117/-1/-1/0/AFD-110324-008.pdf>.*

Stellman, Jeanne Mager, Steven D. Stellman, Richard Christian, Tracy Weber, and Carrie Tomasallo. 2003a. "The extent and patterns of usage of agent orange and other herbicides in Vietnam." *Nature*, 422(6933): 681–687.

Stellman, Jeanne Mager, Steven D. Stellman, Tracy Weber, Carrie Tomasallo, Andrew B. Stellman, and Richard Christian. 2003b. "A geographic information system for characterizing exposure to Agent Orange and other herbicides in Vietnam." *Environmental Health Perspectives*, 111(3): 321–328.

Tveite, Håvard. 2015. "The QGIS Thin Greyscale Image to Skeleton Plugin." <http://plugins.qgis.org/plugins/ThinGreyscale/>.