“Heatwaves and Health at Birth”

Alejandra Paz Rivera-Vicencio*

November 2025

Abstract

This paper examines the impact of prenatal exposure to extreme heat on birth outcomes in Spain. I link over 460,000 administrative birth records (2012–2021) to municipality-level, high-resolution temperature data to construct trimester-specific measures of in-utero heat exposure. Heatwaves are defined as at least three consecutive days above the municipality-specific 95th percentile, with alternative thresholds used for robustness. Unlike most existing studies, I find that heatwave exposure—especially in the third trimester—increases birth weight and gestational age and slightly raises the likelihood of cesarean delivery. The results are stable across multiple specifications and heat definitions. To explore mechanisms, I combine births with new data on local healthcare infrastructure. The positive effects of heatwaves are significantly larger in municipalities with greater access to primary care, suggesting that robust healthcare systems help mothers adapt to extreme heat. Overall, the findings highlight the importance of local resilience in shaping climate-health impacts.

JEL Codes: I10, J13, Q54, R23.

Keywords: Heatwaves, Birth Weight, Gestational Age, Climate Change, Adaptation, Healthcare Access.

*E-mail: alejandrapaz.rivera@ub.edu

1 Introduction

Climate change is increasing the frequency, intensity, and duration of extreme heat events, raising urgent concerns for public health. Heatwaves are particularly concerning for vulnerable populations, including the elderly, outdoor workers, and pregnant women. Fetal development is a critical stage for long-term health and human capital formation (Currie and Rossin-Slater, 2013), and temperature shocks during gestation are often believed to have lasting adverse effects.

A large literature documents that prenatal exposure to heat impairs fetal growth and shortens gestation. Heat stress may lead to dehydration, impaired placental function, or premature labour (Wyrwoll, 2023). Consistent with these mechanisms, Chersich et al. (2020) report strong global associations between high ambient temperatures and low birth weight, preterm birth, and stillbirth. Wang et al. (2013) find that exposure to heatwaves in late pregnancy increases preterm birth risk in Brisbane, while Sun et al. (2020) show that heat, when combined with pollution and limited green space, raises preterm delivery rates in California.

Despite this broad evidence, little is known about how heat exposure affects birth outcomes in Southern Europe—a region projected to experience some of the sharpest increases in heatwave intensity (IPCC, 2023). Moreover, existing studies mostly document adverse effects, leaving open whether and why these patterns might differ across climatic, social, or institutional contexts. Spain provides an especially valuable case to investigate this question. The country combines substantial spatial and temporal variation in temperature with universal access to prenatal and delivery care, offering an environment where the biological effects of heat can be studied with minimal confounding from healthcare access or income disparities (Keivabu and Cozzani, 2022).

This paper examines the impact of in-utero exposure to extreme heat on neonatal outcomes in Spain. I link individual-level birth records from the Spanish Vital Statistics to high-resolution daily temperature data at the municipality level for 2012–2021. Using this merged dataset, I construct trimester-specific measures of exposure to hot days and heatwaves, defined in percentile-based thresholds. Contrary to the prevailing evidence, I find that exposure to heatwaves during pregnancy increases birth weight and gestational age. These effects are concentrated in the third trimester and are robust across multiple definitions of heat exposure, model specifications. The results suggest that, in the Spanish context, short-term heat stress may accelerate delivery only when extreme and sustained, while moderate exposure may act through physiological or behavioural adaptations that favour fetal growth.

A key question raised by these results is why moderate heat exposure appears beneficial rather than harmful. Several mechanisms may help explain this pattern. From a physiological perspective, mild heat stress can increase maternal blood flow and metabolic activity, potentially enhancing nutrient delivery to the fetus. At the behavioural level, heatwaves may prompt pregnant women to reduce physical activity, increase hydration, or spend more time resting indoors, behaviours that could support fetal growth. Institutional factors may also play a role: during periods of extreme heat, prenatal monitoring and medical vigilance may intensify, and air conditioning use may help offset the adverse physiological effects of heat. Later sections of the paper examine these potential pathways by exploring heterogeneity in effects across socioeconomic groups, urban environments, and newborn characteristics.

This paper contributes to several strands of the literature. First, it adds new evidence to the growing body of research on the effects of temperature on health and human capital. While most studies document negative impacts of prenatal heat exposure on birth outcomes, my findings reveal the opposite pattern in a Southern European setting. This challenges the conventional narrative that extreme heat uniformly harms fetal development and suggests that contextual factors—such as adaptation, healthcare access, and climate norms—play a critical role.

Second, the paper contributes to the literature on climate adaptation and health. By documenting a setting where the short-term consequences of heat exposure appear non-adverse, it highlights the potential importance of behavioural and institutional responses—such as hydration, rest, or air conditioning—in mitigating physiological stress. Understanding these mechanisms can inform adaptation strategies in regions facing increasing climate variability.

Finally, this work contributes methodologically by combining administrative birth microdata with high-frequency meteorological information to construct precise, trimester-level measures of in-utero heat exposure. This fine temporal matching allows me to isolate gestational periods most sensitive to temperature shocks, advancing empirical approaches to studying climate–health interactions.

2 Related Literature

Research across economics, epidemiology, and environmental health has increasingly documented that prenatal exposure to extreme temperatures can affect fetal development and neonatal outcomes. A broad consensus links high ambient temperatures to higher risks of low birth weight, preterm delivery, and stillbirth (Chersich et al., 2020; Wang et al., 2013; Chen et al., 2020; Richards et al., 2022). The proposed mechanisms include dehydration, impaired placental function, and inflammation during maternal heat stress

(Wyrwoll, 2023). These effects are consistent with a growing body of evidence showing that heat can act as a physiological stressor that disrupts gestation and accelerates delivery.

Recent work highlights that the interaction of heat with other environmental stressors may amplify health risks. Sun et al. (2020) and Ha et al. (2024) find that concurrent exposure to high temperatures, air pollution, and wildfire smoke significantly raises the likelihood of preterm birth in California. These findings point to the importance of studying compound exposures and to potential synergies between environmental shocks. At the same time, social and economic factors shape both exposure and vulnerability: Cushing et al. (2022) show that low-income and racially marginalised mothers face disproportionate heat-related risks, underscoring the role of adaptation capacity and structural inequality.

Beyond physiology and exposure, institutional and behavioural mechanisms also appear relevant. For instance, Costa-Ramón et al. (2018) show that physician incentives affect cesarean section timing in Spain, while Costa-Ramón et al. (2022) links cesarean delivery to adverse long-term child outcomes. These findings suggest that obstetric practices can mediate the relationship between environmental stress and newborn health. Similarly, Currie and Rossin-Slater (2013) and Quintana-Domeque and Ródenas-Serrano (2017) provide evidence that maternal stress from natural disasters or violence reduces birth weight, supporting the broader notion that environmental shocks influence pregnancy through both physiological and behavioural pathways.

Most of the empirical evidence on heat and birth outcomes comes from the United States, China, and Australia, with relatively little work on Southern Europe. This gap is relevant because the region is projected to experience some of the world's most severe increases in heatwave frequency and intensity (IPCC, 2023). In Spain, Díaz et al. (2018) and Roldán et al. (2016) document rising heat-related mortality and health burdens, while Keivabu and Cozzani (2022) show that high temperatures reduce fertility, particularly in regions with limited access to cooling. However, no study has yet examined how in-utero heat exposure affects birth outcomes using nationally representative administrative microdata.

This paper contributes to filling this gap by providing the first large-scale evidence for Spain on the relationship between gestational heat exposure and newborn health. Unlike most existing work, which finds detrimental effects, I find that heat exposure during pregnancy—especially in the third trimester—increases birth weight and gestational length. This contrasts with prior findings and opens new questions about adaptation, behaviour, and the role of local climate and health systems in shaping the consequences of extreme heat.

3 Data

3.1 Overview

This paper combines rich administrative microdata on births with high-resolution meteorological information to examine how exposure to extreme heat during gestation affects health at birth in Spain. The analysis draws on detailed records from the Spanish Vital Statistics (Estadística de Nacimientos), which provide exhaustive information on all live births and stillbirths, and links them to municipality-level daily temperature data from the E-OBS climate reanalysis dataset. The merged dataset integrates demographic, clinical, and meteorological dimensions at a fine temporal and spatial scale, allowing me to identify exposure to hot days and heatwaves for each pregnancy and estimate their effects on multiple birth outcomes.

3.2 Dependent Variables

The main outcome variable is birth weight (PESON), measured in grams and available for every newborn. I complement this with two additional indicators that capture other dimensions of neonatal health: gestational age at delivery (SEMANAS), expressed in completed weeks of gestation, and a binary indicator for cesarean delivery (CESAREA), distinguishing between vaginal and cesarean births. These outcomes are standard in the health economics and epidemiological literature and enable a multidimensional assessment of how maternal exposure to heat during pregnancy affects both fetal growth and obstetric complications.

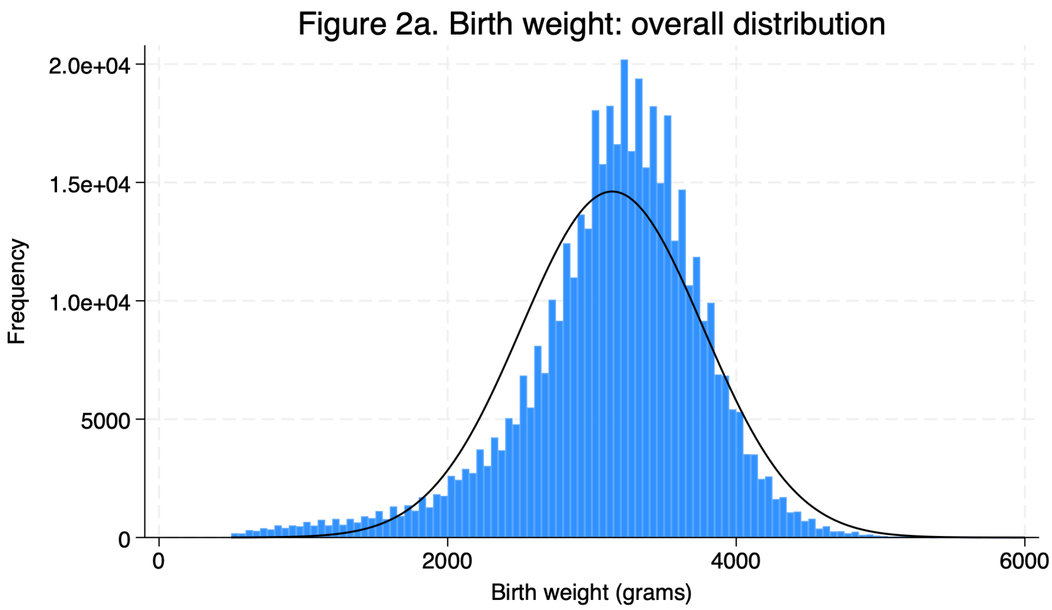

(a) Birth weight: overall distribution

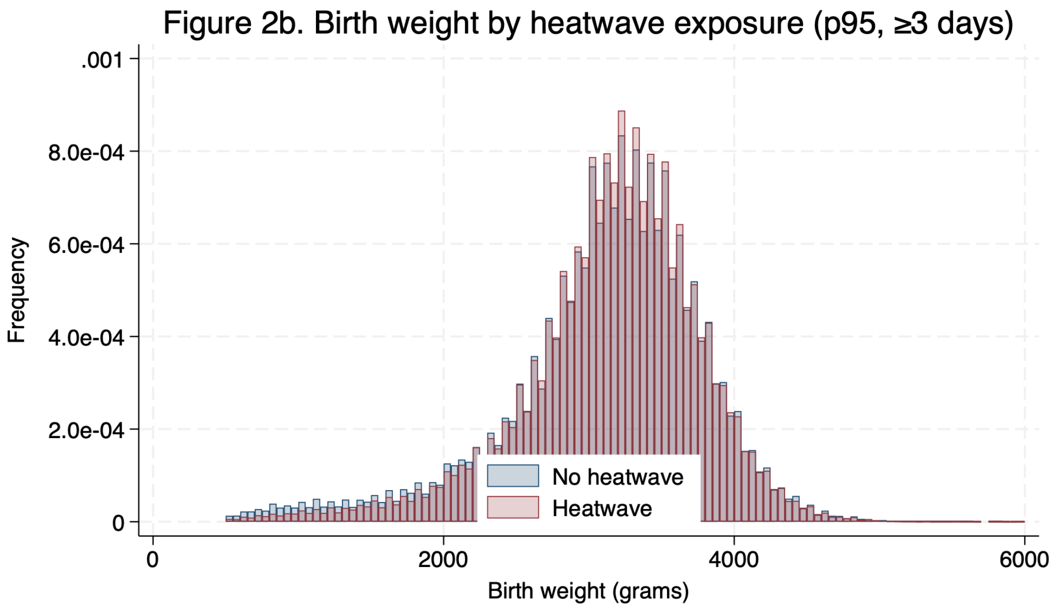

(b) Birth weight by heatwave exposure (p95, $\geq$3 days)

Figure 1: Distribution of birth weight and comparison by heatwave exposure

Notes: Panel (a) shows the overall distribution of birth weight for the full sample of births between 2012 and 2021. Panel (b) compares the birth weight distribution for pregnancies exposed and unexposed to a heatwave ($\geq$3 consecutive days above the municipality-specific 95th percentile of Tmax).

Figure 1 illustrates the distribution of birth weight in the sample. Panel (a) shows that birth weight is approximately normal, centred around 3,250 grams, with a slight right tail. Panel (b) compares distributions for pregnancies exposed and unexposed to heatwaves. The two curves overlap closely but show a modest rightward shift for exposed births, consistent with the regression results that heatwave exposure, particularly late in gestation, increases average birth weight.

3.3 Main Independent Variable: Heat Exposure

To measure exposure to extreme temperatures during gestation, I link each individual birth record to a daily series of maximum temperatures (Tmax) corresponding to the mother's municipality of residence (cod_ine) throughout the entire gestational period. This linkage produces a day-level panel in which each row represents one day of gestation for a specific pregnancy. From this expanded dataset, I compute detailed indicators of heat exposure—such as the number of hot days and the occurrence of heatwaves—based on both absolute and percentile thresholds calculated from local temperature distributions. These measures are then aggregated to the pregnancy level, resulting in one observation per birth that combines demographic and clinical characteristics with comprehensive indicators of temperature exposure. The final dataset covers all births in Spain between 2012 and 2021. To link daily temperature data to each pregnancy, I constructed a unique identification code that distinguishes individual births based on a combination of observable maternal and birth characteristics. This identifier allowed me to expand the dataset to the day level and then collapse it back to the pregnancy level while preserving one-to-one correspondence between records. In a number of cases, identical identifiers were generated for distinct records with the same observable characteristics. In such instances, I retained one randomly selected record to avoid duplication. This procedure ensured internal consistency of the dataset and preserved a representative sample of births for the analysis.

3.4 Sample Construction and Representativeness

I construct a unique birth identifier using [mother's age, parity, marital status, municipality, month–year of birth, sex] and keep the first non-missing record when duplicates arise. This yields $\sim$460,000 births for 2012–2021. I link births to municipality-level daily Tmax from E-OBS using the mother's municipality of residence, aggregate to gestational day, and collapse to trimester-specific exposure measures. Observations with missing municipality codes or implausible gestational ages are dropped.

3.5 Controls and Mechanisms

The birth records include a rich set of maternal and child characteristics that I use both as controls and as potential mechanisms. Control variables include maternal age, education level, and parity (number of previous children), which capture socioeconomic and demographic factors correlated with both fertility and health outcomes. I also control for the sex of the newborn, given evidence that male fetuses are generally more vulnerable to environmental stressors.

To explore mechanisms beyond average birth outcomes, I use indicators of adverse perinatal events. Specifically, the data record whether the infant was born alive (NACVN) and whether the newborn died before or during delivery (MUERN). These outcomes allow me to assess whether extreme heat affects not only fetal growth and gestational duration but also perinatal survival, offering insight into the physiological and behavioural channels through which maternal heat exposure may influence reproductive health.

Appendix Table A1 compares observable maternal characteristics and birth outcomes between pregnancies with and without heatwave exposure. Differences are small in magnitude, confirming that exposed and unexposed groups are broadly comparable along socioeconomic and demographic dimensions.

4 Empirical Strategy

The objective of this paper is to estimate the causal impact of prenatal exposure to extreme heat on health at birth. The main empirical challenge arises because temperature varies systematically with other seasonal or local factors that may also influence pregnancy outcomes—such as diet, physical activity, or disease prevalence. To address these concerns, I exploit high-frequency variation in heatwave occurrence across municipalities and months over time, and estimate fixed effects models of the following form:

$$Y_{impt} = \sum_{j=1}^{3} \beta_j H_{impt}^{(j)} + X'_{impt}\theta + \alpha_p + \gamma_{mt} + \varepsilon_{impt}, \tag{1}$$

where Y_{impt} denotes a health outcome for newborn i in municipality p, born in month m and year t; $H_{impt}^{(j)}$ measures exposure to heatwaves during trimester j of pregnancy; X_{impt} is a vector of individual-level controls; α_p are municipality fixed effects; and γ_{mt} are month-by-year fixed effects. Standard errors are clustered at the municipality level to allow for spatial correlation in exposure to temperature shocks.

Table 1: Table 1. Summary Statistics for Birth, Maternal, and Heat Exposure Variables, 2012–2021

Panel A: Birth outcomes

	Mean	SD	Min	Max	Obs.
Birth weight (grams)	3140.48	631.42	42.00	6500.00	461251
Gestational age (weeks)	38.48	2.64	17.00	46.00	472613
Not Cesarean delivery (share)	0.72	0.45	0.00	1.00	472613

Panel B: Maternal characteristics

	Mean	SD	Min	Max	Obs.
Mother's age (years)	31.98	5.79	12.00	60.00	472613
Number of children in household	1.76	0.91	1.00	14.00	472613
Foreign-born mother (share)	0.30	0.46	0.00	1.00	472613
Married (share)	0.53	0.50	0.00	1.00	472613
Multiple birth (share)	0.03	0.18	0.00	1.00	472613
Male newborn (share)	0.51	0.49	0.00	1.00	245146

Panel C: Heat exposure

	Mean	SD	Min	Max	Obs.
Any heatwave during pregnancy (p95, $\geq$3 days)	0.78	0.42	0.00	1.00	472613
Heatwave in 1st trimester (share)	0.51	0.50	0.00	1.00	472613
Heatwave in 2nd trimester (share)	0.52	0.50	0.00	1.00	472613
Heatwave in 3rd trimester (share)	0.49	0.50	0.00	1.00	472613
Heatwave days during pregnancy (mean)	43.12	124.83	0.00	9240.00	472613
Mean Tmax during pregnancy (°C)	21.39	3.00	8.40	36.05	429636

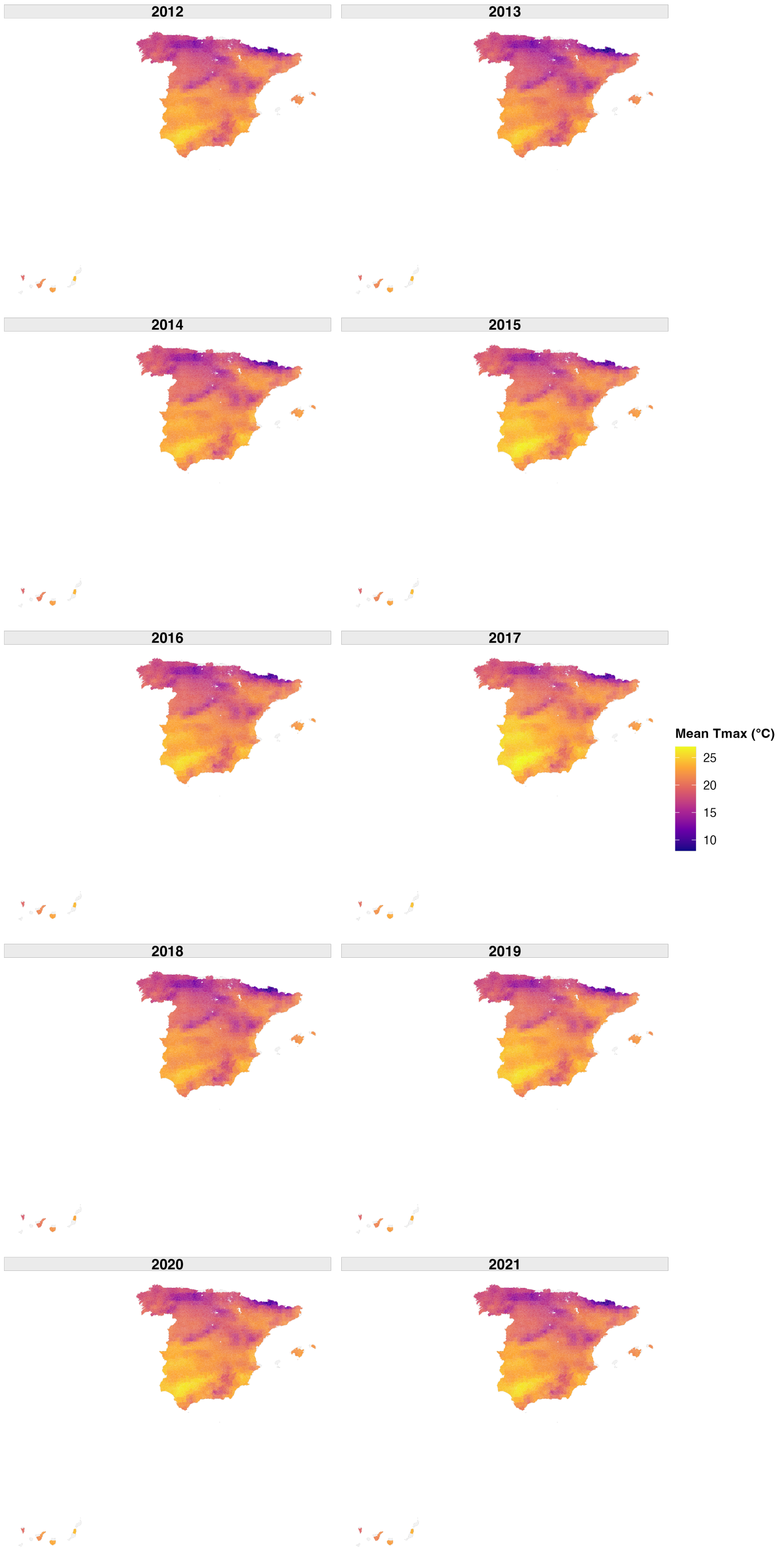

Figure 2: Map of Distribution of Maximum Temperatures in Spain 2012-2021

The model is estimated separately for three outcomes: (i) birth weight (grams), (ii) gestational age (weeks), and (iii) cesarean delivery (binary). For the continuous outcomes, I use OLS, and for the binary outcome, I estimate a linear probability model. The identifying variation comes from short-term deviations in heatwave incidence within municipalities across months and years, after controlling for time-invariant local characteristics (e.g., housing quality, long-run climate) and for nationwide seasonality and time trends through month-by-year fixed effects. As a result, the coefficients β_j capture the average within-municipality effect of additional heatwave exposure during each trimester of gestation.

The main explanatory variable, $H_{impt}^{(j)}$, is defined according to the Spanish Meteorological Agency's concept of a heatwave: a period of at least three consecutive days with daily maximum temperature above the municipality-specific 95th percentile of the historical temperature distribution. I also test alternative definitions based on thresholds at the 90th and 99th percentiles, and on absolute cutoffs (e.g., three or more consecutive days above 35°C). In baseline specifications, $H_{impt}^{(j)}$ is a dummy equal to one if a heatwave occurred during trimester j; I later replace it with the total number of heatwave days to capture dose–response effects. This flexible approach ensures that results are not driven by any particular operational definition of extreme heat.

The vector of controls, X_{impt}, includes maternal age (and age squared), education, parity, marital status, foreign-born status, multiple birth, and the sex of the newborn. These variables absorb demographic and socioeconomic heterogeneity that could correlate with both fertility and health outcomes. Including the newborn's sex also accounts for biological differences in fetal responses to environmental stressors.

Potential threats to identification include unobserved factors correlated with both heat and pregnancy outcomes. For example, regional differences in housing insulation or access to air conditioning could mitigate heat exposure effects, and behavioural responses—such as reduced mobility or increased rest during heatwaves—might alter the relationship between temperature and fetal growth. While such mechanisms may affect the magnitude of the estimated coefficients, the inclusion of municipality and month-by-year fixed effects removes bias from time-invariant housing conditions and from any national seasonal pattern in maternal behaviour. Moreover, the short gestational exposure window limits the scope for reverse causality or endogenous timing of conception in response to anticipated temperature shocks.

Following recent work exploiting quasi-random weather fluctuations to estimate causal effects of temperature on human behaviour and welfare outcomes (e.g., Zivin and Neidell, 2014; Schaller et al., 2023; Cohen and Gonzalez, 2024; Santonja et al., 2025), this empirical strategy uses a high-dimensional fixed effects framework to isolate exogenous short-term variation in local temperatures. Conditional on fixed effects and observable controls, daily and monthly deviations in temperature within municipalities are as-good-as-random with respect to unobserved determinants of birth outcomes. This design therefore purges confounding from both spatial heterogeneity (e.g., persistent climatic, infrastructural, or socioeconomic factors) and common temporal shocks (e.g., nationwide seasonal or policy changes), providing a credible basis for causal interpretation of the estimated coefficients.

To further assess identification, I conduct several specification diagnostics and robustness exercises. First, the assumption that local weather fluctuations are quasi-random conditional on fixed effects and controls is well supported in the meteorological literature. Daily temperature shocks are largely driven by atmospheric variability and are orthogonal to individual-level fertility decisions or local socioeconomic conditions within such short time horizons.

Standard errors are clustered at the municipality level to account for spatial and serial correlation in exposure to temperature shocks, following common practice in weather–health research. As a robustness check.

5 Results

5.1 Main Results

Table 2 presents the main estimates of the effect of heatwave exposure during pregnancy on birth outcomes. Panel A shows that experiencing at least one heatwave—defined as three or more consecutive days above the local 95th percentile of maximum temperature—significantly increases all three outcomes. Birth weight rises by roughly 174 grams, gestational length by about one week, and the probability of cesarean delivery by 3.5 percentage points. These coefficients are statistically significant at the one-percent level and remain stable across specifications, suggesting that prenatal exposure to heat is positively associated with both fetal growth and the likelihood of surgical delivery.

Panel B disaggregates exposure by trimester. The estimates indicate that the effects are concentrated late in pregnancy: a third-trimester heatwave increases birth weight by nearly 199 grams and gestational age by 1.2 weeks, compared with smaller effects of roughly 45 grams and 0.24 weeks for exposure earlier in gestation. The impact on cesarean delivery follows a similar pattern, rising by 3.3 percentage points for late-gestation exposure. This trimester gradient suggests that heat acts primarily through mechanisms operating near the end of pregnancy—such as changes in maternal rest, hydration, or clinical management of delivery—rather than through early-gestation biological impairment.

The magnitude of these estimates is economically and clinically meaningful. A 200-gram increase in birth weight corresponds to roughly one-third of a standard deviation, comparable to the effects of major policy interventions on prenatal health. A one-week increase in gestational length represents about 0.55 of a standard deviation, a large and clinically relevant improvement. Finally, the 3.5–3.8 percentage point rise in cesarean delivery reflects roughly 0.25 of a standard deviation in the cesarean rate. The results are therefore not only statistically robust but also substantial in magnitude, pointing to meaningful changes in maternal or clinical conditions during heat exposure.

Figure 3 visualizes these trimester-specific estimates. The pattern is clear: exposure to heatwaves during the first or second trimester has modest effects, whereas third-trimester exposure leads to a substantial and precisely estimated increase in birth weight. The sharp rise in the third trimester underscores that the timing of exposure is crucial and supports mechanisms operating late in pregnancy, such as increased rest, hydration, or medical supervision.

In contrast to the international evidence, which typically documents adverse effects of high temperatures on fetal development, these findings reveal a positive relationship in the Spanish context. The stability of the estimates across specifications supports the interpretation that moderate heat exposure may improve intrauterine conditions through adaptive behavioural or institutional responses that offset physiological stress. Subsequent sections test the robustness of these results to alternative heatwave definitions and examine heterogeneity patterns that shed light on potential mechanisms.

5.2 Robustness to Alternative Heatwaves Definitions

Tables 3 and 4 test the robustness of the main results to alternative definitions of extreme heat. In Table 3, heatwaves are defined as three or more consecutive days above the 90th percentile of the local temperature distribution, while Table 4 adopts a more restrictive definition based on the 99th percentile. Across both specifications, the estimated effects remain positive, statistically significant, and of comparable magnitude to the baseline results in Table 2, confirming that the findings are not sensitive to the precise operationalisation of heat exposure.

Panel A of Table 3 shows that exposure to a p90 heatwave during pregnancy increases average birth weight by 213 grams, gestational age by 1.25 weeks, and the probability of cesarean delivery by 5.2 percentage points. Panel B again indicates that effects are concentrated in late gestation: third-trimester exposure raises birth weight by 284 grams and gestational length by 1.7 weeks. These coefficients are slightly larger than those obtained under the p95 definition, consistent with the broader and more frequent nature

Table 2: Table 2. Effect of Heatwave Exposure During Pregnancy on Birth Outcomes

	(1) Birth weight	(2) Gestational age	(3) C-section
Panel A: Any heatwave during pregnancy (p95, $\geq$ 3 days)			
Any heatwave during pregnancy	174.093***	1.011***	0.035***
(p95, $\geq$ 3 days)	(15.693)	(0.095)	(0.004)
Observations	460,328	471,587	471,587
Adj. R^2	0.136	0.134	0.072
Panel B: Trimester-specific heatwave exposure (p95, $\geq$ 3 days)			
Heatwave in 1st trimester	44.703***	0.247***	0.016***
	(5.688)	(0.031)	(0.002)
Heatwave in 2nd trimester	46.006***	0.239***	0.015***
	(4.504)	(0.025)	(0.002)
Heatwave in 3rd trimester	198.555***	1.187***	0.033***
	(15.667)	(0.094)	(0.003)
Observations	460,328	471,587	471,587
Adj. R^2	0.154	0.169	0.073

Notes: All regressions include municipality and month×year fixed effects. Controls: maternal age (and squared), parity proxy, education (ESTUDIOM), marital status, foreign-born, multiple birth, and sex of newborn. Standard errors clustered at the municipality level. * p<0.10, ** p<0.05, *** p<0.01.

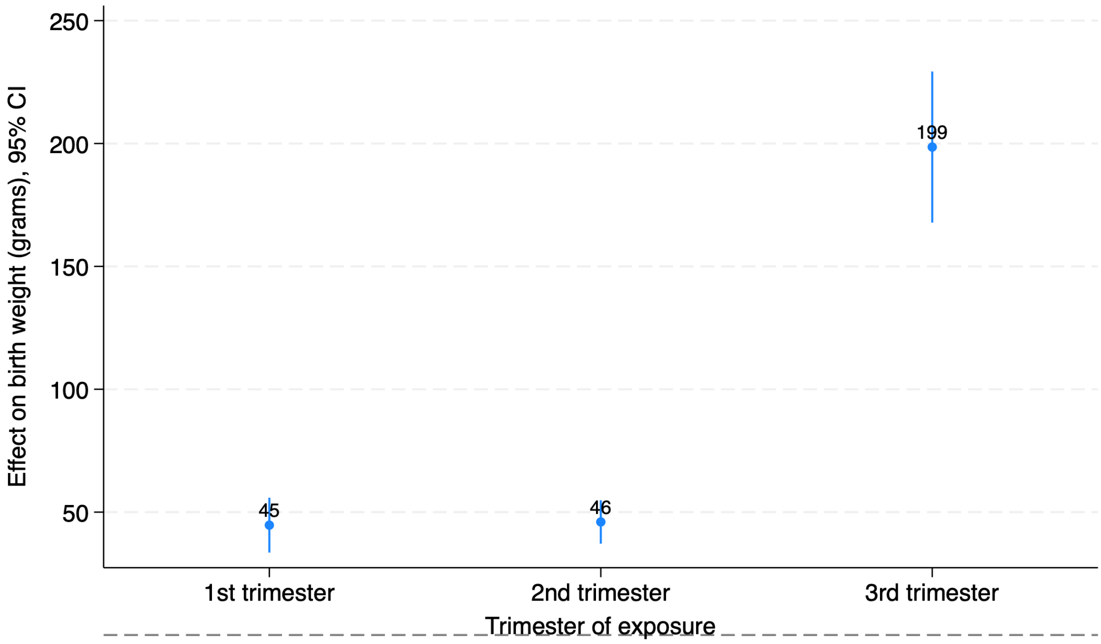

Figure 3: Effect of Trimester-Specific Heatwave Exposure on Birth Weight

Notes: The figure plots point estimates and 95% confidence intervals from the specification in Table 2, Panel B. Each coefficient represents the effect of experiencing at least one heatwave (p95, ≥ 3 days) during the specified trimester of gestation. Vertical bars denote 95% confidence intervals based on standard errors clustered at the municipality level.

of p90 heatwaves, which capture moderate but more sustained exposure.

Results under the p99 definition, reported in Table 4, are remarkably similar. Even when focusing exclusively on the most extreme temperature events, exposure to heatwaves continues to increase both fetal growth and gestational duration. Third-trimester exposure to a p99 heatwave increases birth weight by roughly 136 grams and gestational length by 0.8 weeks, while cesarean deliveries rise by about 2.6 percentage points. The persistence of these effects under increasingly stringent thresholds demonstrates that the positive relationship between heat exposure and neonatal outcomes is not an artefact of a particular cutoff or sample composition, but a robust empirical pattern.

Taken together, these robustness exercises reinforce the credibility of the main findings: independent of how heatwaves are defined, exposure to unusually high temperatures during pregnancy—especially in the final trimester—is consistently associated with heavier and longer-lasting pregnancies. The stability of the estimates across thresholds also mitigates concerns that the results are driven by measurement error or arbitrary classification of extreme heat.

Consistent with a dose–response pattern, Appendix Figure A1 shows that pregnancies exposed to a greater number of heatwave days experience progressively higher birth weights.

5.3 Heterogeneity

Table 5 explores whether the positive effects of prenatal heat exposure on birth weight vary across socioeconomic groups, locations, and by the sex of the newborn. The goal is to assess both heterogeneity in treatment effects and to gain insights into potential mechanisms underlying the main results.

Panel A compares mothers with low versus high education levels. The coefficients are remarkably similar—165 grams and 180 grams, respectively—indicating that the increase in birth weight following heatwave exposure is not driven by differences in maternal education or socioeconomic status. This pattern suggests that the observed effects are unlikely to reflect compositional or behavioural differences linked to income or knowledge and are instead consistent with physiological or contextual adaptations common across groups.

Panel B splits the sample by settlement size. The estimates show that the effect is considerably stronger in large urban areas (+243 g) than in small towns or villages (+137 g). This difference points to environmental and behavioural channels related to heat exposure and mitigation. Urban areas experience more intense and sustained heat due to the urban heat island effect, but also greater access to cooling, medical services, and prenatal mon-

Table 3: Table 3. Effect of Heatwave (p90) Exposure During Pregnancy on Birth Outcomes

	(1) Birth weight	(2) Gestational age	(3) C-section
Panel A: Any heatwave during pregnancy (p90, $\geq$ 3 days)			
Any heatwave during pregnancy	213.364***	1.250***	0.052***
(p90, $\geq$ 3 days)	(50.783)	(0.319)	(0.007)
Observations	460,328	471,587	471,587
Adj. R^2	0.131	0.124	0.072
Panel B: Trimester-specific heatwave exposure (p90, $\geq$ 3 days)			
Heatwave in 1st trimester	36.227***	0.183***	0.014***
	(10.011)	(0.060)	(0.002)
Heatwave in 2nd trimester	38.993***	0.196***	0.015***
	(9.501)	(0.056)	(0.002)
Heatwave in 3rd trimester	283.906***	1.729***	0.044***
	(22.963)	(0.141)	(0.003)
Observations	460,328	471,587	471,587
Adj. R^2	0.162	0.189	0.073

Notes: All regressions include municipality and month×year fixed effects. Controls: maternal age (and squared), parity proxy, education (ESTUDIOM), marital status, foreign-born, multiple birth, and sex of newborn. Standard errors clustered at the municipality level. * p<0.10, ** p<0.05, *** p<0.01.

Table 4: Table 4. Effect of Heatwave (p99) Exposure During Pregnancy on Birth Outcomes

	(1) Birth weight	(2) Gestational age	(3) C-section
Panel A: Any heatwave during pregnancy (p99, ≥ 3 days)			
Any heatwave during pregnancy	158.079***	0.910***	0.036***
(p99, ≥ 3 days)	(12.343)	(0.073)	(0.003)
Observations	460,328	471,587	471,587
Adj. R^2	0.140	0.141	0.072
Panel B: Trimester-specific heatwave exposure (p99, ≥ 3 days)			
Heatwave in 1st trimester	68.491***	0.419***	0.022***
	(5.451)	(0.031)	(0.002)
Heatwave in 2nd trimester	84.430***	0.459***	0.021***
	(6.488)	(0.037)	(0.002)
Heatwave in 3rd trimester	135.918***	0.788***	0.026***
	(10.945)	(0.063)	(0.002)
Observations	460,328	471,587	471,587
Adj. R^2	0.141	0.142	0.073

Notes: All regressions include municipality and month×year fixed effects. Controls: maternal age (and squared), parity proxy, education (ESTUDIOM), marital status, foreign-born, multiple birth, and sex of newborn. Standard errors clustered at the municipality level. * p<0.10, ** p<0.05, *** p<0.01.

itoring. The stronger effects in cities therefor,e suggest that adaptive responses—such as increased rest, air conditioning use, or proactive medical supervision—may amplify the positive association between heat and fetal growth.

Panel C examines heterogeneity by the sex of the newborn. Male and female infants exhibit comparable responses to heatwave exposure, with slightly larger estimates for males (+182 g versus +165 g). This pattern aligns with medical evidence that male fetuses are more sensitive to environmental conditions, potentially displaying stronger growth responses when the intrauterine environment improves or stressors are mitigated. The similarity across sexes nonetheless indicates that the overall mechanism operates broadly rather than being sex-specific.

Overall, these findings reveal limited socioeconomic heterogeneity but some variation across environments, consistent with adaptive behavioural or institutional mechanisms rather than purely biological ones. The larger effects in urban settings in particular suggest that factors such as access to cooling technologies, behavioural adaptation, and clinical management during hot periods may play a key role in shaping how heat exposure affects fetal growth in Spain. At the same time, these heterogeneity patterns should be interpreted as suggestive rather than definitive evidence of mechanisms, as the data do not directly capture behavioural or medical responses during pregnancy.

6 Mechanisms: The Role of Access to Primary Healthcare

A natural question is through which channels heat exposure during pregnancy affects neonatal health. While the medical literature points to pathways related to thermoregulation, dehydration, placental stress, and altered fetal growth, the magnitude of the estimated effects may also depend on the ability of pregnant women to access timely healthcare. In this section, we examine whether local access to healthcare—particularly to Spain's primary care system (Centros de Atención Primaria, CAP)—moderates the effect of heatwaves on birth outcomes.

6.1 Motivation

Primary care plays a central role in Spain's public health system. It is the main point of contact for pregnancy monitoring, early detection of complications, and provision of preventive advice. Access to these services varies across municipalities and provinces due to historical investment patterns and recent reorganisations of CAPs and hospitals.

Table 5: Heterogeneity and Mechanisms

	(1)	(2)
Panel A: By maternal education (Birth weight)		
Columns: Low vs High education		
Any heatwave during pregnancy (p95, $\geq$ 3 days)	165.012***	179.971***
	(15.195)	(18.992)
Observations	142,070	111,130
Adj. R^2	0.125	0.156
Panel B: Large of Cities (Birth weight)		
Columns: Village/Town ($\leq$20k) vs Big City (>100k/capital)		
Any heatwave during pregnancy (p95, $\geq$ 3 days)	136.596***	243.124***
	(17.264)	(20.989)
Observations	171,860	136,695
Adj. R^2	0.128	0.148
Panel C: By sex of newborn (Birth weight)		
Columns: Female vs Male		
Any heatwave during pregnancy (p95, $\geq$ 3 days)	165.724***	182.046***
	(14.861)	(17.051)
Observations	221,500	238,741
Adj. R^2	0.133	0.129

Notes: Each column reports the coefficient on *Any heatwave during pregnancy (p95, $\geq$ 3 days)*. All specifications include municipality and month×year fixed effects. Controls (when available): maternal age (and squared), parity proxy, marital status, foreign-born, multiple birth, sex of newborn, and education dummies. Standard errors clustered at the municipality level. * p<0.10, ** p<0.05, *** p<0.01.

If healthcare access mitigates the physiological and behavioural consequences of heat exposure during pregnancy, we would expect the heatwave effect to be weaker in areas with better primary care availability.

6.2 Data and measurement

Using an administrative census of all healthcare centres in Spain (2009–2024), we construct annual counts of open primary care centres and hospitals at both the municipality and province level. A centre is considered open in a given year if its opening year is at or before that year and its recorded closing year is after that year (if any).

We compute two access measures:

1. the number of primary care centres (CAP) per 10,000 inhabitants;
2. the number of hospitals per 10,000 inhabitants.

For interpretability in interaction regressions, we centre each access variable around its mean so that the main heatwave coefficient reflects the effect at the average level of healthcare access.

6.3 Empirical specification

We estimate the following model:

$$BW_{impy} = \beta_1 HW_{imy} + \beta_2 HC_{my} + \beta_3\, HW_{imy} \times HC_{my} + X_i'\gamma + \lambda_m + \tau_{py} + \varepsilon_{impy}, \quad (2)$$

where HW_{imy} indicates exposure to a heatwave during pregnancy, HC_{my} is the (centred) measure of healthcare access in municipality m and year y, λ_m are municipality fixed effects, τ_{py} are province-by-year fixed effects capturing regional healthcare and demographic trends, and X_i is the vector of maternal controls.

The parameter of interest is β_3, the interaction between heat exposure and healthcare access.

6.4 Results

6.4.1 Municipal access to primary care

We find strong evidence that access to primary care moderates the impact of heatwaves on birth outcomes. Municipalities with more CAPs per capita exhibit higher average birth weight. More importantly, we estimate a negative and statistically significant interaction between heatwave exposure and CAP density:

- the heatwave effect on birth weight is *larger* in municipalities with *low* primary care access;
- it is *smaller*—and can approach zero—in municipalities with *high* primary care access.

This pattern is consistent with health system constraints shaping behavioural or monitoring responses to heat exposure: In areas with limited primary care, pregnant women may adjust their daily routines or receive medical attention later, generating larger changes in fetal growth. By contrast, where primary care access is good, enhanced monitoring and timely advice help buffer these responses, reducing the overall effect.

Table 6 illustrates these heterogeneous effects, showing how the marginal impact of heatwaves declines as municipal CAP availability increases.

6.4.2 Municipal and provincial access to hospitals

In contrast, hospital density does not significantly moderate the heatwave effect at either the municipal or provincial level. This is consistent with the institutional setting: routine prenatal monitoring in Spain is carried out in primary care rather than in hospitals.

6.4.3 Provincial access measures

Provincial-level measures of CAP or hospital availability also do not moderate heatwave effects. This likely reflects the coarseness of provincial aggregates: day-to-day access to pregnancy monitoring is determined locally, not at the provincial scale.

6.5 Interpretation

Overall, these results indicate that *local primary healthcare access is an important mechanism shaping the response of pregnancy outcomes to heat exposure.* Municipalities with fewer primary care centres exhibit stronger heatwave effects on birth weight, while municipalities with better access experience attenuated or null effects. This suggests that prenatal care capacity acts as a protective factor, cushioning the impact of extreme temperatures on maternal health and fetal growth.

Table 6: Mechanisms: Healthcare Access (Municipal Level)

	(1) CAP Density	(2) Hospital Density
Any heatwave (p95, $\geq$3 days)	197.648***	196.386***
	(14.994)	(15.350)
CAP per 10,000 inhabitants (centered)	35.593***	
	(11.352)	
Any heatwave $\times$ CAP (centered)	-38.638***	
	(11.814)	
Hospitals per 10,000 inhabitants (centered)		-37.564
		(38.084)
Any heatwave $\times$ Hospitals (centered)		62.878**
		(31.193)
Observations	424,864	424,864
Adjusted R^2	0.134	0.134

Notes: Dependent variable is birth weight (grams). Centered variables are demeaned at the municipal level. All regressions include municipality fixed effects and province $\times$ year fixed effects. Controls include maternal age (and squared), parity proxy, marital status, education, foreign-born indicator, and multiple birth indicator. Standard errors are clustered at the municipality level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7 Policy Implications and External Validity

The Spanish context—characterised by universal prenatal care, dense primary care coverage, and widespread though unequal access to cooling—appears to mediate the relationship between heat exposure and fetal development. The positive effects on birth weight and gestational length suggest that behavioural adaptations (rest, hydration, changes in activity patterns) and clinical adaptations (more frequent contact with primary care, routine monitoring, and early responses to maternal discomfort) may compensate for physiological heat stress for most pregnancies.

Evidence from the mechanism analysis shows that the heatwave effect is significantly stronger in municipalities with greater availability of primary care centres. This indicates that access to front-line healthcare plays a central role in mitigating heat-related risks during pregnancy. These findings imply three core policy lessons:

1. **Prenatal guidance during heat.** The fact that primary care availability amplifies the positive effect suggests that timely access to routine check-ups, rapid triage, and basic medical advice is valuable during periods of extreme heat. Strengthening communication between health services and pregnant women—hydration reminders, reducing exertion, and scheduling antenatal visits during heat alerts—may be highly effective even without major new investments.

2. **Targeting and healthcare capacity.** Heterogeneity analyses show that adaptation is not uniform. The larger effects in urban areas and in municipalities with more primary care centres indicate that local healthcare capacity shapes the extent to which heatwaves translate into improved (or protected) outcomes. Heat–health action plans could explicitly incorporate pregnancy-specific guidance and ensure that primary care centres have the capacity to respond to increased demand during extreme heat episodes.

3. **External validity and ability to adapt.** The mechanism results underscore that the Spanish pattern—positive effects of heat on fetal development—is likely conditional on access to a robust primary care system. In settings with weaker healthcare infrastructure, limited prenatal monitoring, or constrained access to cooled environments, the balance may shift toward the adverse effects documented in other countries. Understanding which components of adaptation matter most—availability of primary care, emergency response capacity, housing quality, and maternal work conditions—should be a priority for scaling these findings beyond Spain.

In sum, the results indicate that access to healthcare is an important margin of resilience: when primary care is readily available, heatwaves may prompt additional monitoring and behavioural adjustments that protect maternal and fetal health. Where such access is limited, heat exposure may instead exacerbate risks.

8 Conclusion

This paper provides new evidence on the relationship between gestational heat exposure and health at birth, using administrative microdata from Spain linked to high-resolution meteorological records. In contrast to much of the existing literature, which finds that extreme heat impairs fetal development, I show that exposure to heatwaves during pregnancy—particularly in the third trimester—increases both birth weight and gestational length, with modest increases in the probability of cesarean delivery. These results are robust to a wide range of specifications and alternative definitions of extreme heat, and their magnitude is economically and clinically meaningful. Taken together, the findings highlight that the effects of temperature on fetal development are context-dependent and shaped by behavioural, medical, and institutional forms of adaptation.

A further contribution of the paper is to shed light on one mechanism behind these patterns. By merging birth records with newly compiled data on local healthcare infrastructure, I show that the positive effects of heatwaves are substantially larger in municipalities with better access to primary care. This pattern is consistent with the idea that proximity to front-line healthcare facilitates timely monitoring, early detection of discomfort, and rapid medical advice during periods of extreme heat. These results suggest that the capacity of the healthcare system plays a meaningful role in mediating the short-run consequences of heat exposure during pregnancy.

The paper makes three broader contributions. First, it adds large-scale evidence from Southern Europe to the climate–health literature, a region where extreme heat is increasingly common but still understudied. Second, it challenges the prevailing view that high temperatures uniformly harm fetal development, showing instead that in a high-access, highly adapted setting, moderate heat exposure can have neutral or even beneficial effects. Third, it demonstrates the value of combining administrative health data, meteorological information, and institutional features—such as the structure of primary care—to better understand heterogeneity in climate impacts.

At the same time, the study faces limitations that open avenues for future research. The analysis does not observe indoor exposure, air conditioning availability, or maternal time use, nor does it capture detailed information on maternal work conditions or medical leave. These factors likely influence the ability to cope with heat and may interact with access to healthcare. Future work linking administrative health records to data on housing quality, occupational conditions, or local adaptation policies would help clarify how different protective mechanisms operate.

As climate change continues to increase the frequency and intensity of heatwaves, understanding how heat exposure interacts with local adaptation and healthcare systems is essential for designing effective public health responses. The evidence from Spain suggests that when access to primary care is strong, the short-term effects of heat on fetal development need not be negative. Sustaining such resilience in a warming world will require continued attention to both healthcare capacity and the broader set of adaptations that support maternal and fetal well-being.

References

Chen, X., Tan, C. M., Zhang, X., and Zhang, X. (2020). The effects of prenatal exposure to temperature extremes on birth outcomes: the case of china. *Journal of Population Economics*, 33:1263–1302.

Chersich, M. F., Pham, M. D., Areal, A., Haghighi, M. M., Manyuchi, A., Swift, C. P., Wernecke, B., Robinson, M., Hetem, R., Boeckmann, M., and Hajat, S. (2020). Associations between high temperatures in pregnancy and risk of preterm birth, low birth weight, and stillbirths: Systematic review and meta-analysis. *The BMJ*, 371.

Cohen, F. and Gonzalez, F. (2024). Understanding the link between temperature and crime. *American Economic Journal: Economic Policy*, 16(2):480–514.

Costa-Ramón, A., Kortelainen, M., Rodríguez-González, A., and Sääksvuori, L. (2022). The long-run effects of cesarean sections. *Journal of Human Resources*, 57:2048–2085.

Costa-Ramón, A. M., Rodríguez-González, A., Serra-Burriel, M., and Campillo-Artero, C. (2018). It's about time: Cesarean sections and neonatal health. *Journal of Health Economics*, 59:46–59.

Currie, J. and Rossin-Slater, M. (2013). Weathering the storm: Hurricanes and birth outcomes. *Journal of Health Economics*, 32:487–503.

Cushing, L., Morello-Frosch, R., and Hubbard, A. (2022). Extreme heat and its association with social disparities in the risk of spontaneous preterm birth. *Paediatric and Perinatal Epidemiology*, 36:13–22.

Díaz, J., Carmona, R., Mirón, I. J., Luna, M. Y., and Linares, C. (2018). Time trend in the impact of heat waves on daily mortality in spain for a period of over thirty years (1983–2013). *Environment International*, 116:10–17.

Ha, S., Abatzoglou, J. T., Adebiyi, A., Ghimire, S., Martinez, V., Wang, M., and Basu, R. (2024). Impacts of heat and wildfire on preterm birth. *Environmental Research*, 252.

IPCC (2023). Intergovernmental panel on climate change: Climate change 2023: Synthesis report. ipcc, geneva, switzerland. Technical report, Intergovernmental Panel on Climate Change.

Keivabu, R. C. and Cozzani, M. (2022). Extreme heat, birth outcomes, and socioeconomic heterogeneity. *Demography*, 59:1631–1654.

Quintana-Domeque, C. and Ródenas-Serrano, P. (2017). The hidden costs of terrorism: The effects on health at birth. *Journal of Health Economics*, 56:47–60.

Richards, M., Huang, M., Strickland, M. J., Newman, A. J., Warren, J. L., D'Souza, R., Chang, H. H., and Darrow, L. A. (2022). Acute association between heatwaves and stillbirth in six us states. *Environmental Health: A Global Access Science Source*, 21.

Roldán, E., Gómez, M., Pino, M. R., Pórtoles, J., Linares, C., and Díaz, J. (2016). The effect of climate-change-related heat waves on mortality in spain: uncertainties in health on a local scale. *Stochastic Environmental Research and Risk Assessment*, 30:831–839.

Santonja, A., Schmitz, L., and Vall, J. (2025). Understanding the link between heat and intimate partner violence. IEB Working Paper 2025/09, Institut d'Economia de Barcelona (IEB), University of Barcelona, Barcelona. IEB Working Paper 2025/09.

Schaller, M. F. E. L. G. J., thank Ben Hansen, W., Mullins, J., Ortiz-Bobea, A., Park, J. R., Wichman, C., Evans, M. F., Gazze, L., and Robert, J. S. (2023). Nber working paper series temperature and maltreatment of young children. Technical report.

Sun, Y., Ilango, S. D., Schwarz, L., Wang, Q., Chen, J. C., Lawrence, J. M., Wu, J., and Benmarhnia, T. (2020). Examining the joint effects of heatwaves, air pollution, and green space on the risk of preterm birth in california. *Environmental Research Letters*, 15.

Wang, J., Williams, G., Guo, Y., Pan, X., and Tong, S. (2013). Maternal exposure to heatwave and preterm birth in brisbane, australia. *BJOG: An International Journal of Obstetrics and Gynaecology*, 120:1631–1641.

Wyrwoll, C. S. (2023). The heat is on: how does heat exposure cause pregnancy complications?

Zivin, J. G. and Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26.

Appendix

Table A1: Covariate Balance: Pregnancies With vs. Without Heatwave Exposure

	No heatwave	≥1 heatwave	Diff.	p-value
Panel A: Maternal characteristics				
Maternal age (years)	31.73	32.05		0.000
Education: High (share)	0.21	0.25		0.000
Foreign-born (share)	0.32	0.30		0.000
Panel B: Birth outcomes				
Birth weight (grams)	3107.49	3149.88		0.000
Gestational age (weeks)	38.20	38.55		0.000

Notes: Means by exposure group. p-values from two-sided t-tests of equal means. Exposure indicator: any heatwave during pregnancy (p95, ≥3 days).

Table A2: Robustness: Heatwave Definition >35°C (≥3 days)

	(1) Birth weight	(2) Gestational age	(3) C-section
Panel A: Any heatwave >35°C			
Any heatwave >35°C (≥3 days)	160.430***	0.963***	0.030***
	(13.562)	(0.082)	(0.003)
Observations	461,207	472,572	472,572
Adjusted R^2	0.129	0.133	0.071
Panel B: Trimester-specific exposure >35°C			
1st trimester HW (>35°C)	117.475***	0.690***	0.026***
	(9.376)	(0.055)	(0.003)
2nd trimester HW (>35°C)	46.174***	0.276***	0.012***
	(5.815)	(0.029)	(0.003)
3rd trimester HW (>35°C)	192.235***	1.189***	0.034***
	(18.316)	(0.117)	(0.004)

Notes: Each column reports estimates from OLS regressions. Panel A shows the effect of experiencing at least one heatwave above 35°C lasting ≥3 days during pregnancy. Panel B shows trimester-specific effects using the same definition. All regressions include municipality fixed effects and month × year fixed effects. Controls include maternal age (and squared), parity, marital status, education, foreign-born status, and multiple birth. Standard errors clustered at the municipality level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

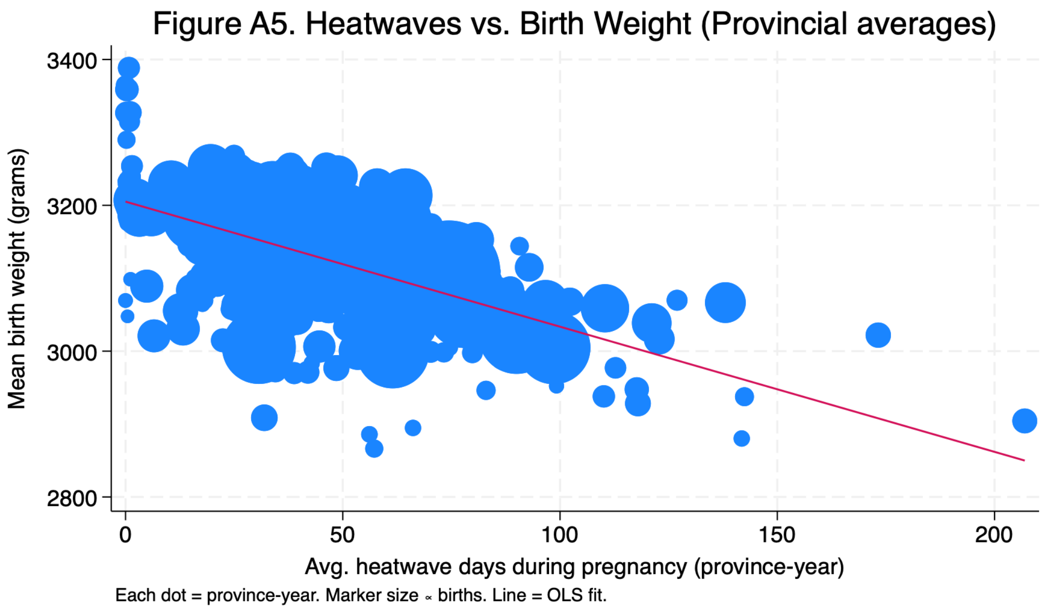

Figure A1: Heatwave and Weight Relation at Province Level

Table A3: Table 6. Mechanisms: Stillbirth and Perinatal Mortality

	(1)	(2)
Panel A: Regression results		
Columns: Stillbirth vs Perinatal death		
Any heatwave during pregnancy (p95, $\geq$ 3 days)	0.000	0.000
	(.)	(0.000)
Observations	472,572	472,572
Adj. R^2	.	-0.001

Notes: Each column reports the coefficient on *Any heatwave during pregnancy (p95, $\geq$3 days)*. The dependent variable is an indicator for stillbirth (col 1) or perinatal death (col 2). All specifications include municipality and month×year fixed effects. Controls: maternal age (and squared), parity proxy, marital status, education, foreign-born, multiple birth, and sex of newborn. Standard errors clustered at the municipality level. * p<0.10, ** p<0.05, *** p<0.01.

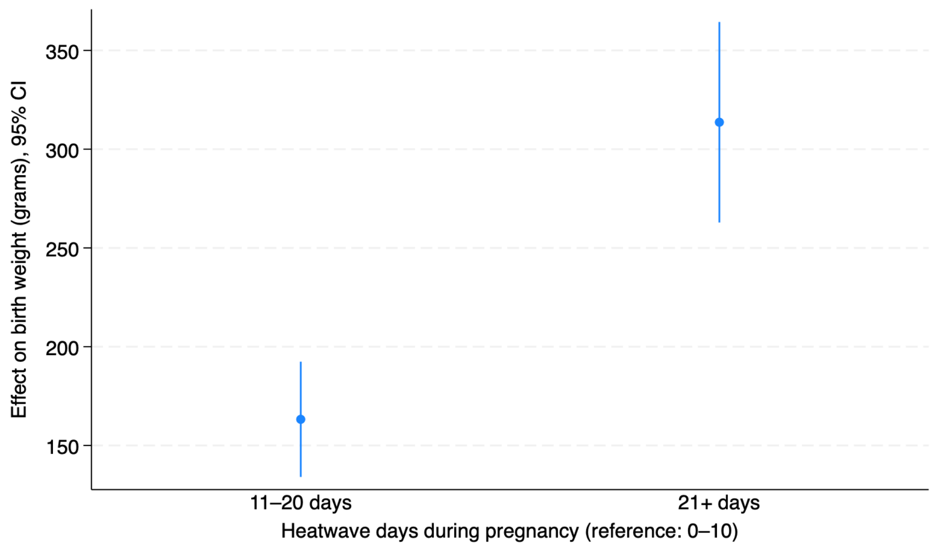

Figure A2: Dose Effect of HEatwaves on Birth Weight