STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
BE551 Degree project in economics
Spring 2026

Make America Fertile Again
What US State Safety Nets Actually Do for Birth Rates

Julia Cook (25949)

Abstract

The United States fertility rate has fallen nearly continuously since 2007, from a peak general fertility rate of 69.3 births per 1,000 women to a provisional record low of 53.1 in 2025 — a decline that standard economic variables explain less than 10% of. This thesis investigates one underexplored candidate: whether state-level social safety net generosity shaped the fertility response to the Great Recession, and through what mechanisms. Using a 51-state panel spanning 2003–2019 and four complementary empirical strategies — a binary difference-in-differences, a continuous DiD, a fixed-effects panel regression, and a triple-differences specification — I find that states entering the recession with more generous safety nets experienced significantly stronger post-2008 fertility trajectories. The average treatment effect on the treated (ATT) is positive and stable across demanding specifications. However, I find little evidence that this advantage operates through the unemployment-fertility cushioning channel: LAMBDA is null across eleven of twelve panel specifications, and the triple-differences result offers only weak support. The event study reveals a fertility differential that builds gradually through the recovery years and stabilizes between 2012 and 2015, which is inconsistent with acute crisis buffering, and more consistent with a structural mechanism. I argue that what safety nets are doing is not replacing lost income in the moment of crisis, but sustaining institutional confidence: the perception, especially among young adults forming families during the recovery, that the consequences of future economic shocks will be collectively rather than individually borne. If that interpretation is correct, the policy implications extend well beyond any single transfer program.

Keywords: Fertility, Welfare, Unemployment, Great Recession, EITC, Difference-in-Differences

JEL: J13, I38, E32, J18, H75, J21

Supervisor: Sampreet Goraya
Date submitted: May 11, 2026
Date examined: May 19, 2026
Discussants: Alexander Ekdahl and Ioana Paduraru
Examiner: Johanna Wallenius

Acknowledgements

I would first and foremost like to thank my family, to whose love and support I attribute all of my academic successes and none of my academic failures. I would like to thank my friends and classmates for their inspiration, as well as my supervisor Sampreet Singh Goraya for his guidance. Finally, I would like to thank everyone who works, in whatever capacity, to make the world a better place for children and the people who raise them.

Contents

List of Tables and Figures

Figures

Tables

1. Introduction

Even if an American wanted to avoid the topic of falling fertility rates, they would have a very hard time doing so. President Trump declared himself the "fertilization president" at a Women's History Month event at the White House (Brooks and Lovelace, 2025). "Our parents aren't having children," stated RFK Jr., the US Secretary of Health and Human Services, at a press conference on expanding fertility treatment access (Kile, 2025). Even the richest man in the world, Elon Musk, will not stop talking, and tweeting, about babies, saying "if people don't have more children, civilization is going to crumble" (Shead, 2021). Interestingly, Musk — who has a net worth of over $800 billion (Forbes, 2025) and served as head of the Department of Government Efficiency (DOGE) — simultaneously supported cuts to federal programs assisting pregnant mothers and children, including a freeze on Head Start grants serving over half a million low–income children (Balingit & Associated Press, 2025).

This statement and tweets point at an interesting tension — why is the birth rate falling? Is it, as the Trump administration has argued, due to a lack of access to fertility care (Kile, 2025)? Or is it that Americans want fewer children on average? Or is it that Americans are choosing not to give birth to children they do not feel they can adequately provide for in an increasingly unstable economy? The fact that stated preferences around intended fertility have remained stable (Guzzo & Hayford, 2023), and that intended conception occurs within one year around ~80% of the time (Gnoth et al., 2003), points at the latter as a key explanation.

The urgency is not merely cultural, but also economic. A falling birth rate carries a fiscal toll: fewer future workers mean fewer future contributors to Social Security and Medicare, compressing the ratio of workers to retirees — a ratio that has already fallen from about four per beneficiary from 1980 to 2008, to roughly three–to–one today, and is projected to decline to 2.2–to–one by 2099 (Board of Trustees, 2025) — and undermining the long–run solvency of the backbone of how America cares for the retired. The Congressional Budget Office's long–run projections already reflect this pressure, with Social Security's trust fund projected to be depleted by 2033 without legislative action (Board of Trustees, 2025).

The numbers are clear about where the drop–off started: the Great Recession. The US fertility rate was trending upward until 2007, where it peaked at a GFR of 69.3 births per 1,000 women aged 15–44 (Centers for Disease Control and Prevention, 2026) before falling almost perfectly continuously from 2008 onward. Interestingly, this trend continued to occur across years when unemployment was near historic lows (Bureau of Labor Statistics, 2026) and real median household income was rising (United States Census Bureau, 2026). This gives leave to a puzzle: what is causing the decline, and how can it be stopped?

This thesis investigates one underexplored part of the answer: whether state–level social safety net generosity shaped the fertility response to the Great Recession, and, if so, through what mechanisms. To investigate this question, I use a 51–state panel spanning 2003–2019, restricted to avoid the confounding effects of the COVID–19 pandemic on both fertility and labour markets and employ four complementary empirical strategies. First, a binary median–split difference–in–differences (DiD) produces an average treatment effect on the treated (ATT) and allows for the creation of an event study, which traces the time path of the two groups. Second, a continuous

DiD design uses each state's pre–recession 2008 safety net generosity as a fixed treatment variable to verify the sign and significance of the ATT (the DELTA estimand). Third, a fixed–effects panel regression tests whether safety net generosity cushions a state's unemployment–fertility relationship (the LAMBDA estimand). Finally, a triple–differences specification adds unemployment as a third dimension to directly test whether the safety net cushions the fertility response to economic shocks.

The results do not answer the puzzle, but they do narrow it down slightly. States that entered the Great Recession with more generous safety nets — as measured by combined TANF, SNAP, federal and state EITC, and ACTC benefits (Schmidt, Shore–Sheppard and Watson, 2024) — experienced significantly stronger post–2008 fertility trajectories. DELTA and ATT are positive and statistically significant, stable even across demanding specifications. Notably, the event study reveals a fertility differential that builds throughout the recovery years and stabilizes between 2012 and 2015, as opposed to peaking at the height of the crisis and then attenuating to zero. However, I find very little evidence that this advantage operates through the unemployment–fertility cushioning effect examined: LAMBDA is null across eleven of twelve panel specifications, and the triple–differences result is only weakly negative. The divergence between these two findings invites the question of what the safety net is actually doing to fertility and through what mechanisms. What the data suggest, then, is not that safety nets are irrelevant to fertility, but that their relevance is perhaps more structural than marginal. States with more generous safety nets produced stronger fertility trajectories not by replacing income in the moment of crisis, but — I argue — by sustaining the kind of institutional confidence that makes parenthood feel viable in the first place. If that interpretation is correct, the policy implications are complex, as it points to a deeper relationship between institutional support and reproductive decision–making that no single program can address.

My thesis proceeds as follows. Section 2 provides background on the Great Recession, the US safety net context, and the trajectory of US fertility since 2007. Section 3 reviews the relevant literature across four key areas — the economics of fertility, economic shocks and fertility, housing and fertility, and social programs and fertility — before addressing possible theories of secular decline. Section 4 describes the methodology and identification strategy, while Section 5 presents the data. Section 6 reports findings, which Section 7 then discusses and interprets. Finally, Section 8 concludes.

2. Background

2.1. The Great Recession: What Happened, to Whom, and for How Long

The Great Recession, according to the NBER, lasted officially for 18 months, December 2007 to June 2009 (NBER, 2010), making it the longest post–WWII recession yet. After a period of unsustainable growth in housing prices, propped up by irresponsible and underregulated lending practices, house prices began to fall in late 2007, triggering the collapse of subprime mortgages. For financial institutions who had invested heavily in mortgage–backed securities, such as Bear Stearns and Lehman Brothers, this collapse spelled disaster: the collapse of these institutions sent shockwaves through credit markets, freezing lending and triggering a broader economic contraction. GDP declined 4.3% from peak to trough (Q4 2007 to Q2 2009), during which period the S&P 500 lost roughly 57% (Federal Reserve History, 2013).

The consequences of the contraction were not relegated to shareholders, as unemployment rose to 10.0% in October 2009, from below 5% in 2007, more than doubling (Bureau of Labor Statistics, 2012). Young workers aged 16–24 were among the hardest hit, with youth unemployment peaking at 19.5% in April 2010 (Bureau of Labor Statistics, 2018). Due to a freeze on hiring, long–term unemployment (27+ weeks) reached historic highs, with the share of unemployed people out of work for six or more months exceeding 45% by 2010 (Bureau of Labor Statistics, 2018). Overall, unemployment was the highest it had been nationally since the 1981–1982 recession, and no state was completely spared (Bureau of Labor Statistics, 2012).

However, the effects of the recession were not completely uniform. Due to the "mancession" nature of the job cuts, with job losses heavily concentrated in construction and manufacturing, state unemployment ranged from over 10% (Nevada, California, Michigan) to 5.2% or lower (North Dakota, Nebraska, South Dakota) at the peak of national unemployment (Bureau of Labor Statistics, 2012). States with large manufacturing and construction industries, such as those in the Rust Belt and Sun Belt, sustained the largest shocks, whereas states whose economies revolve around natural resources were largely spared. The housing bust was also not uniformly spread across states, with "sand states" (Nevada, Florida, Arizona, California) seeing home prices fall the most on average (Olesiuk & Kalser, 2009). The national average was approximately 30% from peak to trough (Federal Reserve History, 2013). Of course, recovery was similarly uneven, with the national unemployment rate not returning to pre–recession levels until 2016, and state–level recovery timelines varying considerably throughout this period (Bureau of Labor Statistics, 2018).

While NBER found the recession's trough to be in June 2009, real median household income did not recover to pre–recession levels nationally until 2016, underlining the magnitude of the recession's effect on American families (US Census Bureau, 2026a). The federal government reacted to the crisis with the ARRA (American Recovery and Reinvestment Act), signed in February 2009, which included tax credits, unemployment insurance, Medicaid support, infrastructure spending, and safety–net expansions, among other measures. At enactment, CBO estimated the total cost at $787 billion over fiscal years 2009–2019, later revised upward to approximately $840 billion (Congressional Budget Office, 2015).

2.2. The US Context and Safety Net

Like most advanced economies, the US offers government support to individuals and families facing economic difficulties. Having a child in the US often involves costs that are not incurred in other countries, such as preschool, which costs on average approximately $9,200 per year nationally for center–based care (Gascon & Hernández Kent, 2025). Combined with the lack of guaranteed paid parental leave — the federal standard under the Family and Medical Leave Act (FMLA) is 12 weeks of unpaid, job–protected leave (U.S. Department of Labor, n.d.) — many parents rely on aid from the government to provide for their families. The core programs assessed in this thesis are TANF (Temporary Assistance for Needy Families), SNAP (Supplemental Nutrition Assistance Program, formerly Food Stamps), EITC (Earned Income Tax Credit), and ACTC (Additional Child Tax Credit).

TANF is a federal block grant program, established in 1996, through which the federal government provides states with funds to offer cash assistance, job training, and support services to low–income families with children for a maximum of 60 months. According to the Office of Family Assistance, "while TANF jurisdictions must meet certain work participation and cost sharing requirements, they have considerable flexibility with TANF funds to implement programs that best serve their distinct communities" (Office of Family Assistance, n.d.). This leads to great cross–state variation, with maximum monthly benefits for a family of three ranging from under $200 to over $900 (Schmidt et al., 2024).

SNAP, similarly, is federally funded but administered at the state level, leading to large variation in eligibility criteria and benefit levels. SNAP works by providing eligible individuals and families with an EBT (Electronic Benefit Transfer) card, which is loaded monthly with funds to use to buy eligible food items at grocery stores, some online retailers, and, depending on local norms, farmers markets (Food and Nutrition Service, n.d.).

The EITC is a federal program that provides a refundable tax credit to working individuals and families. The amount depends on earned income and the number of qualifying children. For the 2026 tax year, the maximum credit for a childless individual is approximately $649, whereas for a family with three or more children it rises to $8,231 (Internal Revenue Service, 2025a). The credit is designed to increase as families earn more at the lowest levels of income, then plateau before phasing out as income rises past certain thresholds. Many states layer their own EITC programs on top of the federal credit, adding a further dimension of cross–state variation (Schmidt et al., 2024).

Finally, the ACTC is the refundable portion of the Child Tax Credit. The Child Tax Credit provides up to $2,200 per child (2026), but as a nonrefundable credit its value cannot exceed a family's tax liability. By contrast, the ACTC is refundable up to $1,700 per child (2026), meaning it can result in a direct cash refund for eligible parents (Internal Revenue Service, 2025b).

2.3. US Fertility: The Post–2008 Decline

Prior to the recession, the US fertility rate was at a peak: the GFR was approximately 69.3 births per 1,000 women aged 15–44, and the TFR was approximately 2.12 — right around replacement level, and the highest rate since 1990 (Martin et al., 2010). Then came the decline. The GFR fell continuously from 2008 onward, reaching 53.8 births per 1,000 women in 2024, with provisional 2025 data suggesting a record low of 53.1 (Osterman et al., 2025; Hamilton et al., 2026). In TFR terms, the rate has steadily crept down to 2024's 1.63, well below the replacement level of 2.1 (World Bank, 2026). Interestingly, the decline continued through a period of strong economic recovery (2014–2019), which has posed quite the puzzle for researchers. The puzzle intensifies when considering that standard economic variables — unemployment, income, housing — account for less than 10% of the sustained post–2008 decline (Kearney, Levine & Pardue, 2022).

Of course, the intensity of the decline was not uniform across states, driven by both demographic composition and differences in norms and economic realities.

Figure 1. US General Fertility Rate, 2003–2019

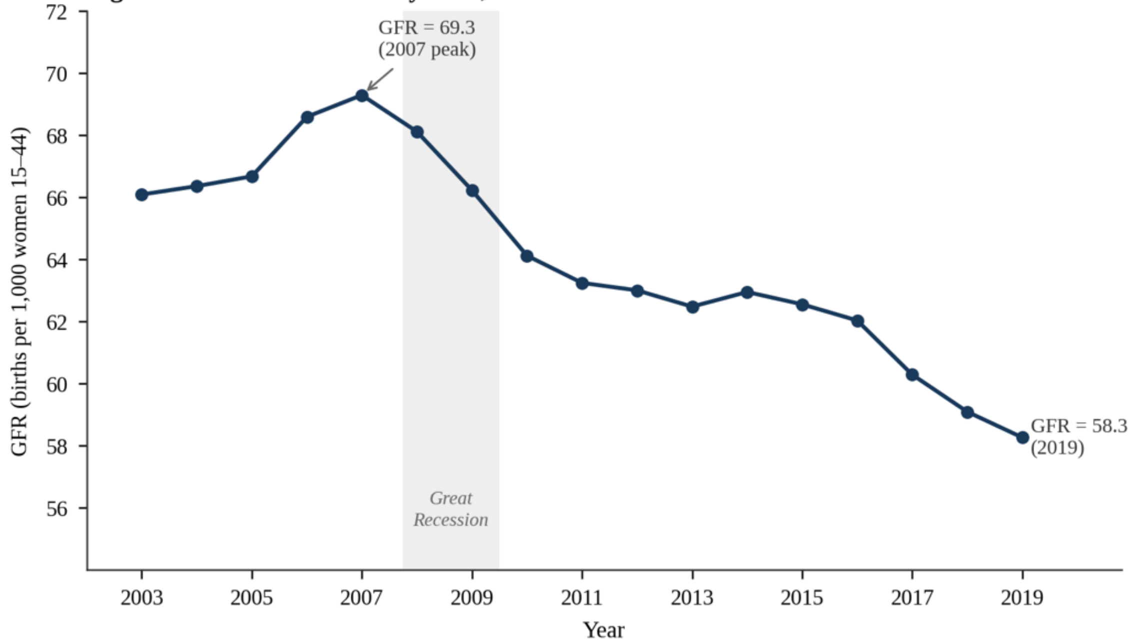

Notes: Population-weighted mean of state-level GFR. Source: CDC WONDER Natality Database.

3. Literature Review

3.1. Economics of Fertility

The modern economic approach to understanding fertility begins with Becker (1960), whose seminal work established that fertility decisions can be understood as rational household choices. In his framework, parents choose how many children to have in order to maximize their lifetime wellbeing, subject to the constraints of their budget. Importantly, Becker distinguishes between the quantity of children and the quality of children, where quality refers to how much parents spend on each child. He argues that as income rises, parents tend to spend more on each child rather than have more children, meaning the income elasticity of quality is larger than that of quantity. Building on this, Becker and Lewis (1973) show that quantity and quality interact: as families have more children, the cost of investing in each one rises, and vice versa. This interaction can produce a negative relationship between income and the number of children even when the pure income effect is positive, which helps explain some of the ambiguity in the income–fertility relationship.

To understand that ambiguity, it is helpful to distinguish between the income effect and the substitution effect. The income effect refers to the increase in demand for a good that comes from having more purchasing power: if children are treated as normal goods, higher income should lead to more children. The substitution effect, on the other hand, refers to a change in demand that results from a change in the relative price of a good. In the context of fertility, the key price is the opportunity cost of the mother's time; specifically, the wages she forgoes by spending time raising children instead of working. When female wages rise, raising children becomes more expensive relative to other uses of time, which tends to reduce fertility. The empirical consensus is that for female wages, the substitution effect dominates the income effect (Mincer, 1963; Butz & Ward, 1979; Galor & Weil, 1996). Male wages, by contrast, increase household income without significantly increasing the opportunity cost of childcare, so the income effect dominates.

Several scholars have highlighted just how central the opportunity cost of female time is to understanding fertility decisions. Schultz (1973), drawing on Mincer (1963), shows that a wife's wage rate is negatively related to fertility because the primary cost of raising children is the time it takes, which is time that falls disproportionately on mothers. Willis (1973) formalizes this in a lifecycle model, arguing that even for mothers who are not currently working, the relevant cost is the reservation wage (the value of the market time they are forgoing). This means the substitution effect applies regardless of whether a mother is actually in the labour force.

With this framework in place, the empirical income–fertility literature becomes easier to navigate. The empirical literature does show a consistent negative correlation between income and fertility across countries and over time (Jones & Tertilt, 2008). However, causal studies that isolate exogenous shocks to male income — specifically, income gains accruing to men in coal–boom regions (Black et al., 2013) or areas experiencing fracking–driven labor demand (Kearney & Wilson, 2018) — tend to find positive effects on fertility. These results are less contradictory than they first appear. The negative cross–sectional correlation reflects the fact that higher–income households face higher opportunity costs of parental time, spend more on each child, and are more likely to live in environments with higher female wages and education, all of which raise the cost of having more children. By contrast, the causal studies isolate an increase in male income that does not change the mother's opportunity cost, which produces a more straightforward positive income effect.

3.2. Economic Shocks and Fertility

One well–established finding in the fertility literature is that fertility is procyclical: birth rates tend to rise during economic expansions and fall during recessions. This pattern was documented long before modern economic theory could explain it (Yule, 1906; Galbraith & Thomas, 1941), and has since been confirmed by over a century of US data and cross–national evidence from OECD countries (Sobotka, Skirbekk & Philipov, 2011). However, an important question remains: do recessions cause people to permanently have fewer children, or do they simply cause people to delay having children? This distinction between *tempo effects* (changes in the timing of births) and *quantum effects* (changes in the total number of children born) is important for understanding how persistent the fertility impact of a recession really is.

Sobotka, Skirbekk, and Philipov (2011) review evidence from across OECD countries and find a broadly negative short–run relationship between rising unemployment and birth rates, while stressing that this relationship is not uniform or mechanically causal. Two of their findings are worth highlighting. First, what seems to matter more than the absolute level of economic hardship is how conditions compare to people's expectations — economic deterioration relative to what people anticipated is a stronger depressant of fertility than a low but stable economic environment. Second, because recessions appear to primarily cause postponement of births rather than permanent decisions not to have children, estimates of the quantum effect based on period fertility rates are likely overstated.

Currie and Schwandt (2014) offer US evidence on this question. Using over 140 million birth records from 1975 to 2010, they track cohorts of women defined by state and year of birth, and relate their fertility to the unemployment rates they experienced at each age. The biggest immediate effect occurs when women are between 20 and 24: a one percentage point increase in the unemployment rate reduces conceptions by approximately 6 per 1,000 women in that five year range. When these cohorts are followed all the way to age 40, however, the cumulative effect rises to 14.2 lost conceptions per 1,000 women, largely driven by women who remain permanently childless. This suggests that at least some of the timing effect is not recovered: recessions do not just delay births, they eliminate some of them. Schaller (2016) adds further nuance by separately estimating the effects of male and female labour demand. She finds that improvements in male labour market conditions raise fertility, while improvements in female conditions have smaller, negative effects, which is consistent with the substitution effect logic discussed in Section 3.1. Interestingly, general unemployment rates, she finds, are associated with only modest decreases in state–level birth rates.

Kearney, Levine, and Pardue (2022) try to explain why US fertility has declined so persistently since 2007. They first establish that the decline is not simply the result of demographic changes, noting that within–group fertility reductions occurred simultaneously across all subgroups. They then test ten standard economic and policy explanations using a state–year panel model, and find that all ten together explain only 6.2% of the total decline. The unemployment channel alone predicts a 3.5% fertility decline during the Great Recession, compared to an observed decline of 7.2%. Additionally, birth rates continued to fall even after unemployment had fully recovered after 2010. The remaining gap, which the authors call the "puzzle," is tentatively attributed to a broad, cohort–level shift in preferences and norms around childbearing (Lesthaeghe & van de Kaa, 1986). Because this kind of cultural shift operates at the cohort level rather than varying year–to–year across states, it is absorbed by year fixed effects in a standard panel model and cannot be separately identified.

3.3. Housing and Fertility

Housing is one of the largest expenses most households face, and when children are involved, space requirements, and therefore housing costs, tend to increase. It is therefore reasonable to expect that housing markets play an important role in fertility decisions. The evidence supports this, and in fact suggests that housing costs may matter more for fertility than broader economic conditions like unemployment.

A key distinction in this literature is between homeowners and renters. Lovenheim and Mumford (2013) exploit variation in housing price growth across metropolitan areas during the late 1990s and early 2000s housing boom to estimate the effect of housing wealth on fertility among homeowners. They find that a $100,000 increase in housing wealth increases total fertility by approximately 0.22 children for women ages 25–44 (which reflects 19% of this cohort's TFR) with no corresponding effect among renters in the same areas. For owners, rising house prices create a wealth effect: the home becomes more valuable, households feel wealthier, and they are more willing and able to have additional children. For renters, that wealth gain does not exist, so rising prices function primarily as a cost increase.

Dettling and Kearney (2014) examine the renter–owner asymmetry more directly. Using MSA–level house prices interacted with local baseline ownership rates, and instrumenting with geographic variation in housing supply elasticity, they confirm that a negative substitution effect applies to all households, while the positive home equity effect only applies to owners. Their most striking finding is a quantitative comparison: at the mean ownership rate, a $10,000 increase in housing prices leads to roughly a 5% increase in fertility for homeowners, compared to a 2.4% decrease for renters; at the mean US ownership rate, this equates to a 0.8% increase in current period fertility rates. To provide context, a one standard deviation increase in the unemployment rate leads to only a 2.1% decrease in births. This comparison suggests that housing market conditions deserve serious attention alongside broader macroeconomic indicators in any empirical analysis of fertility.

3.4. Social Programs and Fertility

Standard economic theory predicts that cash transfers should increase fertility by reducing the effective cost of having children. In practice, the relationship between social programs and fertility is considerably more complicated. Transfer programs are not simple income shocks, as they come with conditions attached, including requirements around family structure, employment, and the number of existing children. Each of these conditions can introduce incentive effects that complicate or offset the basic income effect. Additionally, the assumption that fertility is a planned, income–sensitive decision may not hold for low–income populations, where unintended pregnancy rates are high (Brown & Eisenberg, 1995). As a result, the net fertility effect of any given program depends on how it is designed and who it reaches.

The largest body of work studies Aid to Families with Dependent Children (AFDC), which conditioned eligibility on single parenthood, before ending in 1996. Early studies using variation in benefit levels across states find a positive correlation between benefit generosity and non–marital fertility for white women, but weak or null effects for Black women — a difference the literature never satisfactorily explained (Moffitt, 1998). Moffitt also shows that estimated effects shrink substantially once more state–level controls are included, suggesting that cross–state comparisons may be picking up omitted regional factors, like conservative social norms in Southern states that are correlated with both low benefits and high marriage rates. The key methodological takeaway is that studies comparing benefit levels across states tend to find larger effects than studies tracking changes within states over time.

The 1996 Personal Responsibility and Work Opportunity Reconciliation Act (PRWORA) replaced the AFDC and allowed states to implement family caps, which denied additional benefits for children born while a mother was already receiving assistance. The theoretical logic is straightforward: raising the marginal cost of an additional child should reduce fertility. The evidence does not support this prediction. Kearney (2004) finds that family caps do not reduce birth rates and are even positively correlated with fertility among likely eligible populations. A likely explanation is that financial deterrents have limited effect when unintended fertility dominates: three–quarters of pregnancies among women below the poverty line are unintended (Brown & Eisenberg, 1995).

The EITC is now the largest means–tested cash transfer program in the US: by 1996, federal EITC spending was already 1.7 times federal AFDC spending (Eissa & Liebman, 1996). Because the EITC is conditioned on earned income and the presence of qualifying children, it creates both income and substitution effects on fertility. Baughman and Dickert–Conlin (2009), using a state–year panel with fixed effects, find that increases in the base EITC are negatively correlated with first birth rates for white women, and that the incremental EITC is negatively associated with higher–order births for white women, suggesting the substitution effect dominates. Effects for non–white women are consistently null. This is broadly consistent with Eissa and Liebman (1996), who find that the 1987 EITC expansion increased labour force participation among single mothers by up to 2.8 percentage points overall, and by 6.1 percentage points among those without a high school diploma.

Finally, SNAP differs from the programs above in that it does not condition benefits on family structure, removing the single–parenthood incentive that complicated AFDC–fertility estimates. Almond, Hoynes, and Schanzenbach (2011), exploiting the county–by–county rollout of the Food Stamp Program in the 1960s and 1970s, find that program exposure improved birth outcomes but had no statistically significant effect on fertility rates. The authors interpret the null fertility results as evidence against endogenous sample selection bias in their birth weight estimates, rather than as substantive findings. However, it still is an interesting result. Combined with evidence from Hoynes and Schanzenbach (2009) that recipients treat food stamp benefits as equivalent to unrestricted cash, this null result suggests that income transfers alone, at least at these magnitudes, do not significantly shift fertility decisions.

3.5. Secular Decline

The three economic channels reviewed above — the income and substitution effects, economic uncertainty, and housing costs — help explain short–run variation in fertility. What they do not explain is the sustained, long–run decline in US fertility since 2007. Making sense of that decline requires looking beyond year–to–year economic fluctuations and engaging with two broader frameworks: the second demographic transition and more recent evidence on cohort fertility goals in the US.

Lesthaeghe (1983) describes the second demographic transition (SDT) as a long–run shift characterized by the postponement of marriage and first births, driven by a combination of structural changes (i.e. longer time in education, less stable early career trajectories) and cultural changes, particularly a growing emphasis on individual self–realization and keeping future options open. The period total fertility rate falls during the postponement phase, but whether it recovers depends on whether people have children later to compensate. This recuperation varies widely across countries: in France and Scandinavia, for example, later–age fertility largely makes up for earlier delays, while in German–speaking countries, for example, it does not, and cohort fertility falls as a result.

Beaujouan (2023) challenges the implicit assumption behind the postponement–recuperation framework: that people have stable fertility intentions which they are simply deferring. She synthesizes cohort fertility data and survey research and argues that most young people do not have concrete, long–term fertility goals, and that stated intentions are unstable and weakly predictive of actual behavior. She prefers the term "delay" over "postponement," precisely because it does not carry the implication that catch–up was intended. This matters practically:

conception rates decline sharply after age 32, so women who delay first births face a meaningfully higher risk of involuntary childlessness, regardless of what they had planned.

In the US context, Guzzo and Hayford (2023) use synthetic cohorts from multiple waves of the National Survey of Family Growth to assess whether declining fertility reflects changing preferences or growing difficulty achieving stable goals. They find that intended fertility has remained roughly stable at around two children across cohorts, and that intentions to remain childless rarely exceed 15%. However, more recent cohorts are achieving lower fertility at younger ages than earlier cohorts did, and women in their early forties with fewer children than they may have wanted are increasingly unlikely to report unmet fertility desires, suggesting that goals are quietly revised downward as time passes and options narrow.

Together, these findings point to an important limitation of the empirical literature on fertility. The dominant driver of the long–run US fertility decline is a cohort–level phenomenon involving delay, incomplete recuperation, and the gradual revision of fertility goals that operates independently of the year–to–year economic variation that standard panel models are designed to capture. Year fixed effects absorb this secular trend, making it invisible to the econometrician's standard toolkit. The result is a literature well–equipped to estimate how economic shocks shift fertility at the margin, but less equipped to explain why, over the course of a generation, Americans have come to have fewer children than they say they want.

3.6. Research Gap

The literature establishes that fertility is procyclical; economic shocks suppress birth rates in the short run; housing costs can cut in opposite directions depending on ownership status; and individual social programs have limited effects on fertility (sometimes even negative), with income transfers alone appearing insufficient to move decisions at the margins. Still, none of these mechanisms explains the post–2008 decline: standard economic variables account for less than 10% of it (Kearney et al., 2022), and the dominant force appears to be a slower, cohort–level process of delay and downward goal revision that panel models are structurally unable to identify (Guzzo & Hayford, 2023).

What the literature has not asked is whether the *overall* generosity of a state's safety net — the aggregate signal it sends about what kind of support families can expect — has a measurable effect on how individuals' fertility in those states responds to an economic crisis. This is a different question from whether any single transfer program raises birth rates at the margin. I test this by asking whether states with more generous pre–recession safety nets experienced systematically different post–2008 fertility trajectories, and whether that difference operates through the unemployment–fertility channel, or through some other, undefined mechanism.

4. Methodology

4.1. Difference–in–Differences

I employed two DiD specifications using each state's pre–recession (2008) safety net generosity as the treatment variable. This is plausibly fixed before the onset of the Great Recession (elaborated in 5.2), and therefore plausibly exogenous to fertility dynamics of 2008 and beyond.

4.1a. Primary DiD

The primary analysis involved dividing the states into two groups: states were classified as treated if 2008 cashfood_rs was greater than or equal to the national median, and were otherwise assigned as a control group. The key estimand was the ATT, or the average post–2008 fertility differential between above– and below–median states, relative to pre–2008 trends. There were four specifications: S1 includes only state fixed effects and the treatment interaction, S2 adds trend controls (national unemployment rate and linear time trend), S3 allows each state to have its own linear time trend, and S4 includes the full range of variables (female labor force participation rate, share of women with a bachelor's degree or higher, real median household income). The estimating equation was:

$$ln(GFR)_{it} = \alpha_i + \gamma_t + \delta \cdot (D_i \times Post_t) + X_{it}'\beta + \varepsilon_{it}$$

- α_i = state fixed effect, absorbing all time–invariant cross–state differences in fertility
- γ_t = time controls; national unemployment rate and linear time trend in S1–S2, replaced by state–specific linear trends in S3–S4 following Wolfers (2006)
- D_i = 1 if state i's 2008 cashfood_rs was at or above the national median; 0 otherwise
- $Post_t = 1$ if $t \geq 2008$; 0 otherwise
- $X_{it}'\beta$ = vector of time–varying controls in S4: log FHFA HPI, female BA share, and log real median household income, all lagged one year
- δ = ATT: the average post–2008 fertility differential between above– and below–median safety net states, net of state fixed effects and time controls Standard errors clustered at the state level to account for serial correlation within states over time
- $\delta > 0$: above–median states maintained relatively higher fertility post–2008 than below–median states

S3 and S4 replace the common linear time trend with a state–specific linear trend for each of the 51 states, following Wolfers (2006), who employs an analogous specification in his DiD analysis of divorce law and domestic violence as a robustness check on differential pre–period trajectories. The two–group treatment structure remains identical across all four specifications — S3 and S4 still estimate the post–2008 fertility differential between above– and below–median states. What changes is what is controlled for: rather than assuming all states share a common long–run fertility path, the model absorbs each state's individual trajectory before estimating the differential between groups. If the ATT were driven by pre–existing state–level divergence rather than a genuine post–recession effect, it would fall under this specification.

All time–varying control variables in S4 are lagged one year to reflect the approximately nine–month conception window between economic conditions and recorded births.

The primary purpose of this analysis was to provide event study visualization of treatment, as well as to complement DELTA.

4.1b. Secondary DiD

The secondary specification, which was a DiD with a continuous treatment variable, had the demeaned 2008 cashfood_rs value for each state (in thousands of 2022 dollars, relative to the national mean). All 51 states (50 states plus the District of Columbia) were included and the full distribution of generosity was preserved. The estimand of interest, DELTA, theoretically answers the question: for each additional $1,000 of pre–recession safety net generosity, how much better was the post–2008 fertility trajectory? In line with Acemoglu & Finkelstein's 2008 paper on Medicare share in hospitals and the 1983 PPS reform, DELTA will not be interpreted as a structural marginal effect in this thesis (Acemoglu & Finkelstein, 2008).

The general form is:

$$ln(GFR)_{it} = \alpha_i + \gamma_t + DELTA \cdot (SN_{i,2008} \times Post_t) + X_{it}'\beta + \varepsilon_{it}$$

- α_i = state fixed effect, absorbing all time–invariant cross–state differences in fertility
- γ_t = time controls: national unemployment rate and linear time trend
- $SN_{i,2008}$ = demeaned 2008 cashfood_rs value (COLA–adjusted, in thousands of 2022 dollars)
- $Post_t = 1$ if $t \geq 2008$; 0 otherwise
- $X_{it}'\beta$ = vector of time–varying controls in C3: log FHFA HPI, female BA share, and log real median household income, all lagged one year
- DELTA = primary estimand: the post–2008 fertility differential per additional $1,000 of pre–recession safety net generosity; DELTA > 0 indicates more generous states maintained better post–2008 fertility trajectories
- Standard errors clustered at the state level to account for serial correlation within states over time

4.2. Fixed Effects Panel Regression

The next test is a fixed effects panel regression, beginning with a standard within–estimator to estimate the average effect of safety net generosity on fertility, as well as the cushioning effect of safety net generosity on the relationship between unemployment and fertility, holding state and time–trend–specific factors steady. In order to isolate the effect of the specific variables I wanted to investigate, state fixed effects were necessary in order to control for time–invariant characteristics specific to each state, such as culture, geography, history, institutions, etc.

While state fixed effects were an obvious choice, the treatment of time was a more difficult design choice. While the standard approach in the literature is to include year fixed effects, which absorb all common year–to–year variation, year fixed effects combined with state fixed effects and a slow moving policy environment would eliminate much of the variation needed to best identify LAMBDA. The cushioning interaction requires year–to–year cross–state variation in unemployment to be present in the data. By instead including the national unemployment rate directly as a control, alongside a time trend, I remove the common national shock from the residual while leaving cross–state deviations intact.

This is a deliberate identification choice with known trade–offs: in panel designs where the key interaction requires cross–sectional variation in the shock variable to remain in the residual, year fixed–effects undermine the validity of the estimand. This follows Kearney (2004), who estimates a state–panel fertility model with a time–varying welfare policy treatment using state–specific linear time trends and a lagged state unemployment rate in lieu of year fixed effects, which would absorb identifying variation in the treatment variable. Instead, as mentioned, I included the national FRED unemployment rate and a linear time trend to absorb the common shocks and the secular fertility decline without eliminating needed variation. Of course, as mentioned, this choice has tradeoffs: with three time–varying variables controlling for overlapping variation (the national unemployment rate, the linear trend, and the lagged state unemployment rate), the coefficient on the national rate is subject to collinearity and should not be interpreted as a standalone causal estimate. A6 serves as a useful specification check, replacing the national rate and linear trend with year fixed effects.

As specified, the regression is:

$$ln(GFR)_{st} = \alpha_s + \beta \cdot ln(Unemp_{s,t-1}) + \gamma \cdot ln(Policy_{s,t-1}) + LAMBDA \cdot [ln(Unemp_{s,t-1}) \times ln(Policy_{s,t-1})] + \delta \cdot NatlUnempt + \theta \cdot t + \varepsilon_{st}$$

Where:

- α_s = state fixed effect, absorbing all time–invariant cross–state differences in fertility
- β = elasticity of fertility with respect to lagged state unemployment
- γ = elasticity of fertility with respect to lagged safety net generosity (log cashfood, lagged one year)
- LAMBDA = primary cushioning estimand: the differential fertility response to unemployment across levels of safety net generosity (both variables were demeaned before multiplication, so LAMBDA is interpretable at the mean of unemployment and safety net generosity)
- δ = coefficient on national unemployment rate, controlling for common cyclical shocks; not interpreted as a standalone causal estimate due to collinearity with the linear trend and lagged state unemployment
- θ = coefficient on linear time trend, capturing the secular fertility decline
- $ln(Policy_{s,t-1})$ = log of cashfood (unadjusted 2003–2007, COLA–adjusted cashfood_rs 2008–2019), lagged one year
- LAMBDA < 0: cushioning — fertility response to unemployment smaller in more generous states
- LAMBDA = 0: no cushioning
- Standard errors clustered at the state level to account for serial correlation within states over time

While this offered an obvious complement to my analysis, the within–state year–to–year safety net variation is small and potentially endogenous, as states may expand safety nets in response to recessions, which would result in LAMBDA being biased towards zero. Because of this theoretical limitation, this analysis was a supplement to my DiD models.

4.3. Parallel Trends Assumption

The key identifying assumption underlying both of the DiDs, as well as the following event study, is that, absent the differential in pre–recession safety net generosity, above– and below–median states would have followed similar fertility trends (Roth et al., 2023). Violation of this assumption would bias outcome coefficients toward spurious findings.

I formally test the assumption for the binary median–split specification by running an F–test on the four pre–period (2003–2006) treatment–time interaction coefficients. Although this test is not directly applicable to the continuous DiD, both specifications use the same pre–determined 2008 treatment variable and the same panel structure, meaning that the binary pre–trend test provides indirect support for the parallel trends assumption underlying the continuous difference in differences model as well. However, since the F–test is not directly applicable to the continuous DiD, and no analogous formal test is conducted for DELTA, the parallel trends assumption for that specification therefore rests on theoretical rather than empirical grounds.

However, there is an important caveat to the F–test: with only four pre–period years, statistical power is limited — failure to reject parallel trends is not the same as proof that they hold, and the test may lack the sensitivity to detect subtle or non–linear differential trends between the two groups. The event study provides a visual complement to the F–test, allowing the examination of the pre–period coefficients. In addition, controlling for individual state level trends is included as a robustness check, which, as shown in the results, supports the parallel trends assumption.

4.4. Event Study

This was a dynamic extension of the binary DiD. By replacing the single post–recession indicator with a vector of event–time dummies relative to the 2007 reference year (with no additional time controls included, as the dummies absorb common shocks), I was able to visualize the time path of δ_g, the estimated difference in log fertility between above–median and below–median safety net states (note: still using 2008 cashfood_rs to assign groups— explained in 5.2) in that year (after removing state fixed effects). This allows the understanding of the trend of the effect, if any, after the recession.

$$\ln(GFR)_{it} = \alpha_i + \Sigma\_g\ \delta\_g \cdot (D_i \times \mathbb{1}[t-2007=g]) + \varepsilon_{it}$$

- α_i = state fixed effect, absorbing all time–invariant cross–state differences in fertility
- D_i = 1 if state i's 2008 cashfood_rs was at or above the national median; 0 otherwise
- g = event time relative to 2007; g = 0 is the omitted reference category, normalising both groups to zero in 2007
- δ_g = estimated log fertility differential between above– and below–median states at event time g, after removing state fixed effects

- $\delta_g \neq 0$ for $g < 0$ would indicate pre–existing differential trends — a violation of parallel trends
- $\delta_g > 0$ for $g > 0$ indicates the treatment group maintained higher fertility relative to the control group, compared to their common 2007 baseline
- The time path of δ_g across post–period years reveals whether any differential emerged immediately or accumulated gradually

One limitation of the event study interpretation is that treatment assignment is fixed at 2008, meaning the post–period trajectory reflects both the pre–determined structural position of states and any ongoing effects of safety net generosity that correlate with 2008 levels — the two cannot be cleanly separated with this design.

4.5. Triple–differences

The triple–difference specification extends the binary DiD by adding a continuous shock variable, lagged state unemployment, as a third dimension. Where the binary DiD estimates the average post–2008 fertility differential between treatment and control states, the triple difference asks a more specific question: does safety net generosity cushion the fertility *response to unemployment* in the post–recession period? The estimand, LAMBDA_unemp, is the triple interaction of unemployment × treatment group × post–2008 indicator. A negative LAMBDA_unemp would indicate that above–median safety net states had a smaller fertility response to unemployment shocks after 2008 than below–median states (i.e. that the safety net buffered the unemployment–fertility channel specifically, rather than simply shifting the fertility level). The pre–period treatment–control unemployment slope difference provides an analogous pre–trend check for this specification: if the two groups responded differently to unemployment before the recession, the post–period triple interaction would be confounded. An analogous specification replacing unemployment with the FHFA House Price Index as the shock variable is included to test whether the cushioning effect, if any, extends to housing cost shocks.

Unemployment DDD:

$$\ln(GFR)_{it} = \alpha_i + \beta\cdot Unemp_{i,t-1} + \varphi\cdot(Unemp_{i,t-1} \times D_i) + \rho\cdot(Unemp_{i,t-1} \times Post_t) + LAMBDA\_unemp\cdot(Unemp_{i,t-1} \times D_i \times Post_t) + \Sigma\_g\ \delta\_g\cdot(D_i \times \mathbb{1}[t-2007=g]) + \varepsilon_{it}$$

Housing DDD:

$$\ln(GFR)_{it} = \alpha_i + \beta\cdot \ln(HPI_{i,t-1}) + \varphi\cdot(\ln(HPI_{i,t-1}) \times D_i) + \rho\cdot(\ln(HPI_{i,t-1}) \times Post_t) + LAMBDA\_housing\cdot(\ln(HPI_{i,t-1}) \times D_i \times Post_t) + \Sigma\_g\ \delta\_g\cdot(D_i \times \mathbb{1}[t-2007=g]) + \varepsilon_{it}$$

- α_i = state fixed effect
- D_i = 1 if state i's 2008 cashfood_rs was at or above the national median; 0 otherwise
- $Post_t$ = 1 if t ≥ 2008; 0 otherwise
- β = baseline fertility response to the shock in control states pre–2008
- φ = pre–2008 treatment–control difference in shock sensitivity; φ ≠ 0 would indicate differential pre–trends in shock response
- ρ = change in control group shock sensitivity post–2008

- LAMBDA = cushioning estimand: the triple interaction of shock × treatment × post–2008; LAMBDA < 0 indicates treatment states had a smaller fertility response to the shock post–2008
- Both shock variables demeaned before constructing interactions, so LAMBDA is interpretable at the mean of the shock
- $\Sigma\_g\ \delta\_g \cdot (D_i \times \mathbb{1}[t-2007=g])$ = vector of event–time interactions absorbing differential pre–existing fertility trends between groups; g = 0 (2007) omitted
- Standard errors clustered at the state level

5. Data

For all data used in the main analysis, the sample is restricted to 2003–2019 to avoid confounding effects of the COVID–19 pandemic on both fertility and labour markets. All specifications — primary panel regressions, DiD, and event study — use a single unified balanced panel of 51 states × 17 years = 867 state–year observations. The safety net variable uses COLA–adjusted cashfood_rs for 2008–2019 and unadjusted cashfood for 2003–2007, spliced at 2008 to maintain a continuous series across the full panel; the stability of the unadjusted measure across this boundary is documented below in Section 5.2. The DiD treatment variable uses only the 2008 COLA–adjusted cashfood_rs value as a fixed, pre–determined measure of state safety net generosity at recession onset.

5.1. Outcome Variable: General Fertility Rate

The data on general fertility rates was obtained from the CDC WONDER Natality database, via two separate extracts (2003–2006 and 2007–2019) (Centers for Disease Control and Prevention, 2026). This was necessary because the underlying data systems differ before and after 2007. The unit was the general fertility rate, or GFR, which is defined as births per 1,000 women aged 15–44, by state and year. The panel covers all 50 states, as well as DC, for a total of 51 "states." The years in the panel are 2003–2019, and N = 867 (51 states times 17 years) state–year observations. In specifications that include one–year lagged controls (which are all primary and preferred specifications) N = 816 to reflect the dropped lag year of 2003. The outcome is entered in natural log in all regressions.

5.2. Primary Treatment Variable: Safety Net Generosity

The data on safety net generosity comes from Brookings Institution Safety Net Index v2.1 (Schmidt et al., 2024). The index is a state–year panel, which I loaded directly in Stata from Brookings' provided .dta file in order to preserve full decimal precision. The key variable is cashfood, or cashfood_rs (a cost of living (COLA)–adjusted version of cashfood, deflated using BEA Regional Price Parities, which only began in 2008). Cashfood combines TANF, federal EITC, state EITC, ACTC, and SNAP. It is calculated, according to Brookings, using TAXSIM and the Urban Institute Welfare Rules Database. The fact that the COLA–adjusted version begins only in 2008, the first full year of the recession, posed issues for my analysis. This is because, for the continuous and standard DiD, I need to establish a pre–shock baseline as 1) a continuous treatment variable and 2) a group–assignment tool. However, the unadjusted measure is highly stable year–on–year; there is a mean absolute rank shift of 0.9 positions between 2007 and 2008, and a maximum shift of 6 positions. Because the recession did not peak until late 2008, leaving little time for meaningful policy shifts to respond to early recession indicators, the 2007 and 2008 unadjusted rankings mirror each other so closely. In addition, the American Recovery and Reinvestment Act of 2009 was the first meaningful national response to the recession in Cashfood terms. This can also be visualized by looking at the raw cashfood numbers from 2007–2009: it is evident that they remain roughly unchanged 2007–2008, and spike notably in 2009 (see the Appendix). For this reason, I used 2008's COLA measure to assign groups and as the continuous treatment variable.

For the primary panel regressions, which require a continuous safety net variable across the full 2003–2019 panel, I use unadjusted cashfood for 2003–2007 and cashfood_rs for 2008–2019. This introduces a measurement discontinuity at 2008, as the series shifts from nominal to real purchasing power at the treatment year. I considered backcasting COLA–adjusted values for 2003–2007 using each state's 2008 COLA ratio, but rejected this approach after finding that the ratio correlates negatively with nominal benefit levels ($r = -0.38$) — meaning the backcast systematically reordered states by cost of living rather than generosity, contaminating the pre–2008 variation. The null result for LAMBDA is robust to this choice: robustness specifications R11 (restricting to 2008–2019 only, using cashfood_rs exclusively) produces a null LAMBDA (−0.002), confirming the splice does not drive the primary finding.

5.3. Economic Shock Variable: Unemployment

For the main economic shock variable, I use the annual unemployment rate. For the state–level rates, I use BLS LAUS which provides annual rates for all 51 states from 2003–2019 (BLS, 2026). For the national unemployment rate I use Federal Reserve Economic Data's UNRATE, an annual average rate representing "unemployed as a percentage of the labor force. Labor force data are restricted to people 16 years of age and older, who currently reside in 1 of the 50 states or the District of Columbia, who do not reside in institutions (e.g., penal and mental facilities, homes for the aged), and who are not on active duty in the Armed Forces" (Federal Reserve Bank of Saint Louis, 2026). The state unemployment is entered as a natural log, lagged one year to reflect the conception window. The national unemployment is entered in levels, and not as a log, as it functions as a time control: it must absorb the common macroeconomic environment so that the effect of state–level unemployment is not confounded by the national business cycle.

5.4. Housing Cost Variable: FHFA House Price Index

I used the Federal Housing Finance Agency (FHFA) All–Transactions House Price Index (HPI) at the state–year level (Federal Housing Finance Agency, 2026). This provides panel data covering all states, including DC, across the full 2003–2019 panel with no gaps. For the key variable, I used "HPI with 2000 base," which indexes each state's cumulative price appreciation to 100 in the year 2000. This means the variable is not cross–state comparable in levels — a value of ~150 in New York and ~190 in Nevada in 2007, for example, reflect different rates of appreciation since 2000 and not absolute price levels (so New York's ~150 could be higher in absolute prices than Nevada's ~190).

The inability to directly compare levels was not a concern for my analysis. Because all specifications include state fixed effects, the estimator removes each state's mean and exploits only within–state year–to–year variation. In my analysis, I ask versions of: in years when a state's HPI is above its own historical mean, is fertility above or below its historical mean? Therefore, cross–state level differences are never directly compared. This makes the 2000–base index perfectly suited for the needs of this thesis.

5.5. Control Variables

5.5a. Female labor force participation (FLFP):

I source the FLFP data from BLS CPS Table 14, which includes employment status of the civilian noninstitutional population by sex, race, Hispanic or Latino ethnicity, marital status, and detailed age (Bureau of Labor Statistics, 2026). I downloaded separate tables for each year (2003 – 2019) and isolated the group code 03, or "women" (ages 16+). The table includes a "percentage of population" engaged in the "civilian labor force." This percentage was used as opposed to the raw figures, as it is a more relevant control when states vary greatly in size. FLFP is excluded from the primary specification due to potential endogeneity, as women's labor force participation and fertility are often jointly determined, and including FLFP risks absorbing part of the mechanism of interest. It appeals only in robustness specification R10, where it does not materially change the primary findings.

5.5b. Median household income:

For an income control, I use the CENSUS SAIPE, which has "Median Household Income" estimates for all states and years (U.S. Census Bureau, 2026). I then deflate it to real 2022 dollars using BLS CPI–U annual averages (annual average indices were manually transcribed from BLS into the .do file, with a base value of 295.6 for 2022). Because the fixed effects estimator absorbs common time variation, the choice of the base year does not affect the regression estimates.

5.5c. Female education attainment:

Since ACS began in 2005, I backcast using linear extrapolation from each state's 2005–2007 trend to create values for 2003 and 2004. This approach follows the linear extrapolation precedent established in panel data studies that require pre–period observations for variables with delayed data availability.

The underlying microdata comes from IPUMS USA, which delivers ACS extracts as a fixed–width .dat file accompanied by a Stata dictionary file (.do) that reads and labels the data (Ruggles et al., 2025). I built the extract on the IPUMS website selecting the variables sex, age, educd, statefip, perwt, and gq, covering survey years 2005–2019. After loading the data via the IPUMS–provided do–file, I restricted the sample to women aged 25 and over in the civilian non–institutional population, dropped territories and military–identified state codes, and collapsed to state–year shares using person weights (perwt). The resulting variable — the weighted share of women 25+ with a bachelor's degree or higher — was then backcasted as described above to cover the full 2003–2019 panel.

5.6. Descriptive Statistics

The table below presents summary statistics for the key variables across the full 867 state–year observation panel.

Table 1. Descriptive Statistics (N = 867 state-year observations, 2003-2019)

Variable	Mean	SD	Min	Max
General Fertility Rate (births per 1,000 women 15–44)	64.5	7.5	46.8	95.7
State unemployment rate (%)	5.7	2.1	2.1	13.5
Safety net generosity ($ thousands, 2022 USD)	6.109	0.836	4.432	9.227
FHFA HPI (base year 2000)	146.8	31.4	83.6	346.7
Female LFPR (%)	59.6	4.5	47.7	70.1
Median household income (real 2022 USD, thousands)	70.1	11.5	47.9	104.5
Female BA share (% women 25+ with bachelor's or higher)	29.1	6.4	15.7	59.9

Notes: Safety net generosity is the Brookings Safety Net Index cashfood_series (Schmidt, Shore-Sheppard & Watson, 2024), spliced from unadjusted cashfood (2003–2007) and COLA-adjusted cashfood_rs (2008–2019). Median household income deflated to real 2022 USD using BLS CPI-U. Female BA share backcasted to 2003–2004 using linear extrapolation. All variables enter primary specifications as one-year lags.

Sources: Centers for Disease Control and Prevention (2026); Schmidt, Shore-Sheppard and Watson (2024); Bureau of Labor Statistics (2026); Federal Housing Finance Agency (2026); U.S. Census Bureau (2026); Ruggles et al. (2025).

The outcome variable, the General Fertility Rate, averages 64.5 births per 1,000 women aged 15–44, with a standard deviation of 7.5, a minimum of 46.8, and a maximum of 95.7. There is both geographic and time variation encapsulated in these numbers: the minimum is likely Vermont in the late 2010s, whereas the maximum is likely Utah in the early 2000s.

Safety net generosity, the treatment variable, has a mean of $6,109 with a standard deviation of $836. The minimum was $4,432, whereas the maximum was $9,227; the ratio of minimum to maximum is approximately 3x. This variation provides the source of cross–sectional identifying variation for both the LAMBDA and DELTA estimands.

Similarly, the state unemployment rate varies greatly. The mean is 5.7% with a standard deviation of 2.1 percentage points. The minimum is 2.1% and the maximum is 13.5%, a spread of 11.4 percentage points. Again, there is both geographic and time variation in these figures. The within–panel standard deviation is large relative to the mean, which reflects the severity and uneven geographic distribution of the 2008 shock, among other underlying differences between states.

Female labor force participation has an average of 59.6% and standard deviation of 4.5 percentage points. The minimum is 47.7% and the maximum is 70.1%. Most of the variation is cross–sectional rather than within–state over time. Because FLFP and fertility are jointly determined (women tend to reduce labor force participation when they have children) controlling for FLFP risks absorbing some of the very mechanisms I am trying to investigate. Therefore, it is used only as a robustness control.

The median household income (in real 2022 USD) has a mean of $70,057, with a standard deviation of $11,530. The minimum was $47,866 (Mississippi), where the maximum was $104,500 (Maryland or New Jersey). The ratio is slightly more than 2x, and the relatively large standard deviation reflects structural economic differences across states. While these differences are absorbed by state fixed effects, I included income as a control to account for some state's faster post–recession recovery time, reflected in faster increasing real income post–recession. Without controlling for income, it is impossible to isolate the effect of safety net generosity on the outcome variable.

Female educational attainment averages 29.1%, with a standard deviation of 6.4 percentage points, a minimum of 15.7%, and a maximum of 59.9%. The wide range reflects both structural differences across states and the secular increase in female educational attainment over the study period.

The housing cost index, as measured by the FHFA HPI with base year 2000, has a mean of 146.8, and a standard deviation of 31.4. the minimum of 83.6 reflects a state whose within–state price appreciation had reversed below its own 2000 baseline by the trough of the housing bust, whereas the maximum of 346.7 reflects the impressive growth in some other housing markets.

6. Results

6.1. Parallel Trends

In order to proceed with the DiD, as well as present the results, it is crucial to validate the parallel trends assumption. Using a formal F–test on the four pre–period treatment–time interaction coefficients (2003–2006, omitting 2007 to ensure a clean F–test), I found $F(4, 50) = 0.394$ ($p=.812$). The pre–period treatment–control fertility differential is jointly indistinguishable from zero, which supports the assumption of parallel fertility trends for above– and below–median safety net states before the onset of the recession.

Figure 2. GFR by Safety Net Generosity Group, 2003–2019

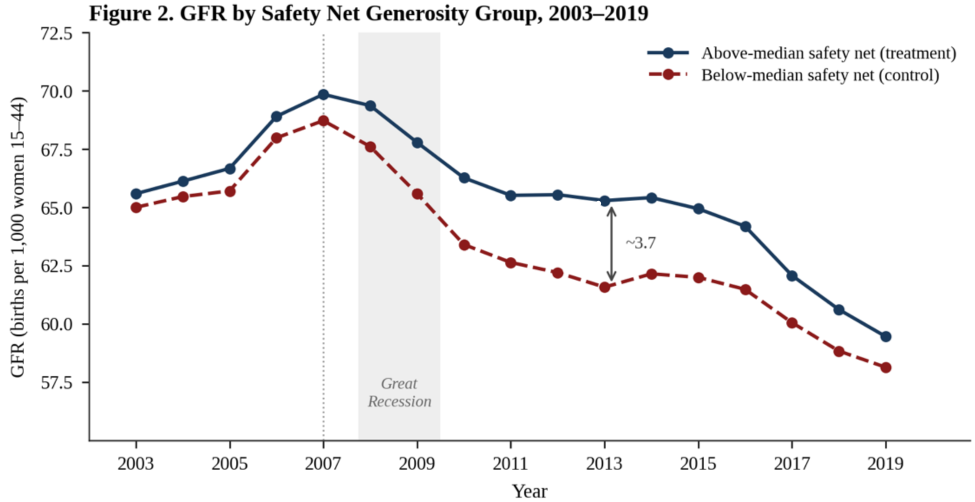

Notes: Unweighted state means within each group. Treatment = above-median 2008 safety net generosity (26 states); control = below-median (25 states). Vertical dotted line marks 2007 reference year.
Source: CDC WONDER Natality Database; Schmidt, Shore-Sheppard & Watson (2024).

6.2. Primary Estimand (ATT)

In Table 2, we see the results for the binary median–split DiD. The 26 above–median states are treated, and the 25 below–median states act as controls. The estimand is the average post–2008 fertility differential between groups, relative to pre–recession trends. For S1 (basic, state FEs and treatment interaction only), the ATT = 0.027*. For S2, which adds the national unemployment rate and linear secular fertility trend, ATT remains unchanged at 0.027*. In S3, which replaces the common national trend with state–specific linear trends for all 51 states, following Wolfers (2006), ATT increases to 0.030***. In S4, with state–specific trends and full controls (HPI, female BA share, household income), ATT = 0.034***. The fact that the estimate strengthens under more demanding specifications is encouraging: if the positive result was driven by pre–existing state–level trends, for example, absorbing individual state trends would attenuate the result, not increase it.

Table 2. Binary DiD: Average Treatment Effect on the Treated (ATT)

Dependent variable: log general fertility rate. Treatment = above-median 2008 safety net generosity (26 states); control = below-median (25 states). All specifications include state fixed effects and clustered standard errors. S3 is the preferred specification.

	(S1) *Basic*	(S2) *+ Time controls*	(S3) *+ State trends*	(S4) *+ Full controls*
ATT: Treatment group × post-2008	**0.027***	**0.027***	**0.030*****	**0.034*****
	(0.0153)	*(0.0153)*	*(0.0097)*	*(0.0081)*
N	867	867	867	816
R^2 within	0.217	0.479	0.749	0.834
Time controls	No	Yes	Yes	Yes
State-specific linear trends	No	No	Yes	Yes
Full controls (HPI, education, income)	No	No	No	Yes
Pre-trend $F(4,50) = 0.394$ [p = 0.812, supports parallel trends]				

Notes: Standard errors in parentheses. ATT is the post-2008 fertility differential between above- and below-median safety net states, relative to pre-2008 trends. Time controls = national unemployment rate and linear time trend. State-specific linear trends follow Wolfers (2006). Full controls add log FHFA HPI, female BA share, and log real median household income, all lagged one year. State-specific trend coefficients (51 states, S3–S4) suppressed for brevity; available on request. Pre-trend F-test: joint significance of four pre-period treatment-time interactions; failure to reject supports parallel trends assumption.

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Source: Author's calculations using CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026).

6.3. Secondary Estimand (DELTA)

Table 3 reports the results of a continuous DiD, where the treatment variable is each state's demeaned 2008 cashfood_rs value. DELTA answers the question: do states with greater pre–recession safety net generosity show systematically better post–2008 fertility trajectories, relative to their own pre–recession baseline? The preferred specification, C2, includes the national unemployment rate and a linear trend, as estimated DELTA =0.035 (SE = 0.010, p <0.01). The estimate remains stable across specifications: DELTA = 0.035*** in the basic specification (C1) and 0.032*** when housing prices, educational attainment, and household income are added (C3). Like ATT, the consistency across increasingly demanding specifications suggests that DELTA is not picking up on correlates of pre–recession state policy generosity. This serves to reinforce the findings of the ATT, and, following Acemoglu and Finkelstein (2008), the coefficient should not be interpreted as a precise causal dose–response, and rather as an indicator of the direction and strength of the relationship.

Table 3. Continuous DiD: Effect of Pre-Recession Safety Net Generosity on Post-2008 Fertility (DELTA)

Dependent variable: log general fertility rate. All specifications include state fixed effects and standard errors clustered at the state level. C2 is the preferred specification.

	(C1) *Basic*	(C2) *+ Time controls*	(C3) *+ Full controls*
DELTA: Safety net generosity × post-2008	**0.035*** **	**0.035*** **	**0.032*** **
	(0.0101)	*(0.0101)*	*(0.0089)*
N	867	867	816
R^2 within	0.240	0.502	0.591
Time controls	No	Yes	Yes
Full controls (HPI, education, income)	No	No	Yes

Notes: Standard errors in parentheses. Treatment variable is each state's demeaned 2008 COLA-adjusted safety net generosity (cashfood_rs, Brookings Safety Net Index v2.1). DELTA > 0 indicates more generous states maintained better post-2008 fertility trajectories relative to their own pre-recession baseline. Time controls = national unemployment rate (FRED UNRATE annual average) and linear time trend. Full controls add log FHFA HPI, female BA share, and log real median household income, all lagged one year. N = 816 in C3 due to one-year lag.
** $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$*

Source: Author's calculations using CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026); Federal Housing Finance Agency (2026); Ruggles et al. (2025); U.S. Census Bureau (2026).

6.4. Event Study

The event study replaces the single post–recession indicator with a vector of year dummies interacted with the treatment indicator, relative to the reference year. The pre–period coefficients are near zero, aligning with the DiD F test. In the post–period, the differential opens gradually, and there is not a sharp jump at 2008 or 2009. The gap stabilizes at approximately 0.030–0.035 from 2012 until it begins narrowing 2017–2019. The gradual development is interesting: it is inconsistent with acute crisis buffering (which would peak early on and then attenuate), instead evoking medium–run structural advantages, which accumulate over the recovery.

Figure 3. Event Study: Dynamic ATT on Log Fertility Rate

Above- vs. below-median pre-recession safety net generosity

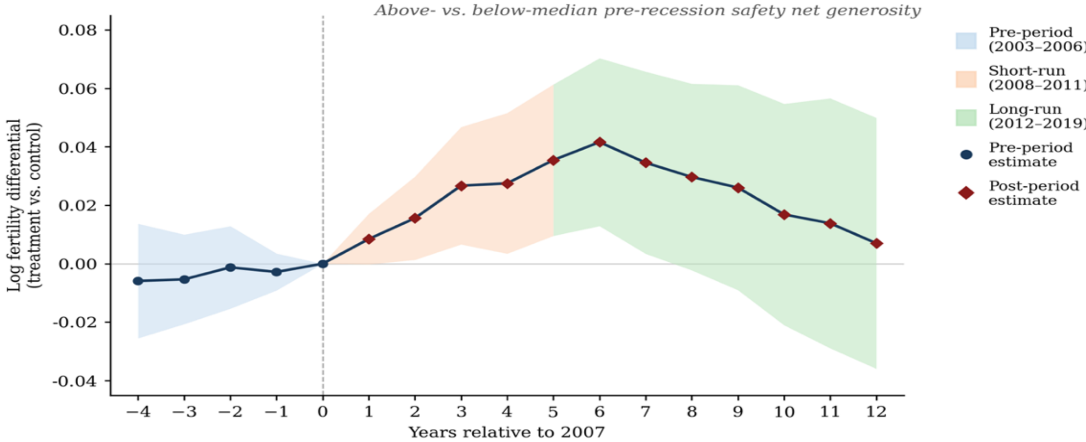

Notes: Shaded bands = 95% confidence intervals. 2007 is the omitted reference year (normalised to zero). Clustered standard errors at state level. All 51 states. Pre-trend F(4,50) = 0.394, p = 0.812.

Source: Author's calculations. CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024).

6.5. Fixed Effects Panel Regressions

Table 4. Fixed Effects Panel Regression: Unemployment Cushioning Estimand (LAMBDA)

Dependent variable: log general fertility rate. All specifications include state fixed effects and standard errors clustered at the state level. N = 816 throughout (one year lost to lagging). A3 is the primary specification.

	(A1) *Unemp only*	(A2) *+ Safety net*	(A3) *+ Interaction (primary)*	(A4) *+ HPI*	(A5) *+ All controls*	(A6) *Year FEs*
Log state unemployment rate (lagged)	-0.052***	-0.050***				
	(0.011)	*(0.014)*				
Log safety net generosity (lagged)		-0.022				
		(0.047)				
Log state unemployment rate, demeaned (lagged)†			-0.051***	-0.030**	-0.030**	-0.080***
			(0.014)	*(0.015)*	*(0.015)*	*(0.026)*
Log safety net generosity, demeaned (lagged)†			-0.017	-0.019	-0.021	0.041
			(0.048)	*(0.046)*	*(0.047)*	*(0.063)*
LAMBDA: Unemployment × safety net (demeaned)			**0.058**	**0.059**	**0.059**	**-0.030**
			(0.064)	*(0.062)*	*(0.060)*	*(0.093)*
Log FHFA House Price Index (lagged)				0.058*	0.058*	
				(0.034)	*(0.030)*	
Log real median household income (lagged)					0.016	
					(0.039)	
Female BA share, % (lagged)					-0.004	
					(0.003)	
National unemployment rate (FRED UNRATE)	0.006***	0.007***	0.006***	0.005***	0.004***	
	(0.001)	*(0.001)*	*(0.001)*	*(0.001)*	*(0.001)*	
Linear time trend	-0.008***	-0.008***	-0.008***	-0.009***	-0.007***	
	(0.001)	*(0.001)*	*(0.001)*	*(0.001)*	*(0.002)*	
N	816	816	816	816	816	816
R^2 within	0.551	0.551	0.552	0.560	0.564	0.687
LAMBDA interaction	No	No	Yes	Yes	Yes	Yes
HPI control	No	No	No	Yes	Yes	No
Full controls	No	No	No	No	Yes	No
Year fixed effects	No	No	No	No	No	Yes

Notes: Standard errors in parentheses. LAMBDA is the interaction of demeaned log state unemployment and demeaned log safety net generosity; both demeaned at panel means before multiplication, making LAMBDA interpretable at the mean of both variables. LAMBDA < 0 indicates cushioning (safety net reduces fertility sensitivity to unemployment). † Demeaned versions replace raw logged variables from A3 onward. National unemployment rate and linear time trend are time controls only; not causal. A6 replaces both with year fixed effects; year FE coefficients suppressed.

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Source: Author's calculations using CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); BLS LAUS (Bureau of Labor Statistics, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026); Federal Housing Finance Agency (2026); Ruggles et al. (2025); U.S. Census Bureau (2026)

Table 4 presents the six FE specifications (N=816 throughout, due to the loss of one year due to lags). All models include state fixed effects and cluster standard errors at the state level.

The most basic model (A1) establishes the core relationship between unemployment and fertility, including only state FEs, lagged log state unemployment, national unemployment, and the linear time trend. The log state unemployment rate is −0.052***. Adding the safety net level in A2 results in a minute decrease to −0.050***, suggesting that the unemployment–fertility relationship is not much affected by safety net generosity levels. This is not surprising given that within–state year–to–year variation in generosity is small and potentially endogenous to economic conditions, which makes the panel FE coefficient an imperfect basis for causal inference.

The key estimand of the panel analysis, LAMBDA (the interaction between demeaned log state unemployment and demeaned log safety net generosity), is introduced in A3. LAMBDA is estimated at 0.058 and is statistically indistinguishable from zero. It remains at 0.059 in A4 and A5 as housing price controls and the full control battery are added. The sign is consistently positive across A3–A5, although insignificant. It is a stable null.

The national unemployment rate enters positively and significantly across A1–A5 (0.004*** to 0.007***). As discussed in the methodology, this coefficient should not be interpreted causally due to collinearity with the linear trend and state unemployment. The linear time trend is consistently negative (−0.007*** to −0.009***), reflecting the ongoing secular fertility decline documented throughout the study period — the same decline that Kearney, Levine and Pardue (2022) note is largely unexplained by standard economic variables. Housing prices enter positively in A4 and A5 (0.058* in both), consistent with the housing wealth effect documented by Lovenheim and Mumford (2013), in which rising home values support fertility among owners through a wealth channel that outweighs the counteracting substitution effect at the margin.

A6 replaces the national unemployment rate and linear trend with year fixed effects, serving as a specification check on the identification approach described in the methodology. Year fixed effects sharpen the state unemployment coefficient to −0.080***, as they absorb national cyclical variation more completely than the combination of linear trend and national unemployment rate. LAMBDA shifts to −0.030 and remains statistically indistinguishable from zero, with a standard error of 0.093 (the largest standard error of the specifications by a large margin: the next largest is 0.064). This sign change reflects a difference in identifying variation rather than a substantive finding: year fixed effects absorb all common time variation more completely than the linear trend, leaving only within–state deviations net of the full national shock, and with less cross–state variation available to identify the interaction, the point estimate turns insignificantly negative. The consistency of the null across both approaches is itself informative: LAMBDA's zero is not dependent on time control choices, but holds regardless of how national variation is handled.

6.6. Triple Difference

Table 5. Triple-Difference: Does Safety Net Generosity Cushion Fertility from Economic Shocks?

Dependent variable: log general fertility rate. Two shock variables tested side by side: state unemployment rate (DDD-1) and FHFA House Price Index (DDD-2). Treatment = above-median 2008 safety net generosity.

	(DDD-1) *Shock: unemployment*	(DDD-2) *Shock: house prices (HPI)*
Baseline shock effect (control, pre-2008)	-0.031***	0.021
	(0.0102)	*(0.0973)*
Shock × treatment group (pre-period slope diff)	0.003	0.158
	(0.0145)	*(0.1447)*
Shock × post-2008	0.040***	-0.134**
	(0.0088)	*(0.0639)*
LAMBDA: Shock × treatment × post-2008 (cushioning)	**-0.021***	**-0.064**
	(0.0127)	*(0.0937)*
N	816	816
R^2 within	0.426	0.387

Notes: Standard errors in parentheses, clustered at the state level. LAMBDA is the triple interaction of the shock variable, treatment group, and post-2008 indicator. LAMBDA < 0 indicates treatment states experienced a smaller fertility response to the shock post-2008 (cushioning). DDD-1 unemployment shock is lagged one year; DDD-2 HPI shock is lagged one year. Both specifications include event-time dummies interacted with treatment group to control for differential pre-period trends. Pre-period treatment-control slope difference = 0.003 (p > 0.10), consistent with parallel pre-trends in the DDD framework.

* p < 0.10 ** p < 0.05 *** p < 0.01

Source: Author's calculations using CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); BLS LAUS (Bureau of Labor Statistics, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); Federal Housing Finance Agency (2026).

To examine whether safety net generosity cushions the response between unemployment and fertility specifically, I used a triple–difference model, which takes the binary DiD design and adds the dimension of unemployment. The estimand, LAMBDA_unemp, is the triple interaction of unemployment x treatment x post–2008. The pre–period treatment–control slope difference was 0.003, statistically indistinguishable from zero. LAMBDA_unemp = –0.021 *, meaning treatment states had a modestly smaller fertility response to unemployment post–2008; however, due to the lack of high–level statistical significance, I therefore do not attempt to decompose the share of the ATT attributable to this channel, as this is insufficient to confirm the cushioning mechanism. An analogous triple–difference specification replacing unemployment with the FHFA House Price Index as the shock variable produces LAMBDA_housing = −0.064, which is a null result, following with the unemployment DDD and suggesting that safety net generosity does not differentially cushion the fertility response to housing cost shocks either.

6.7. Robustness

The full robustness results are reported in the appendix. Across eleven out of twelve panel specifications, LAMBDA is null and consistently positive–signed, confirming that the absence of a cushioning effect is not an artifact of any particular time control, sample restriction, or variable construction choice. DELTA is positive and significant at 0.032–0.035*** across all three DiD specifications, remaining stable whether the basic specification, trend controls, or the full control battery is used.

7. Discussion

7.1. Why LAMBDA Is Null: The Limits of Income Replacement

The divergence between the strong DiD results and the null LAMBDA estimates may provide a new perspective on why states with more social support see more resilient fertility: the mechanisms through which unemployment affects fertility may not be ones that can be "cushioned" by short–run partial income replacement. While existing literature consistently finds that unemployment depresses fertility (Schneider 2015; Schaller, 2016), the assumption that unemployment leads to lower household income which in turn leads to decreased fertility may be overly simplistic. Unemployment changes more than just income: searching for work consumes time and attention, as well as generating stress and making it difficult to plan for the future (will the unemployed person need to move? Start a new type of work? Go back to school?). In this framework, cash and food transfers, while offsetting some of the financial losses, do not address the time and energy components. As Schaller points out, fertility decisions respond not only to household income, but also to opportunity costs and time allocation within households. Unemployment may therefore reduce fertility largely due to uncertainty and time costs, meaning that the safety net generosity may be structurally limited in its ability to "cushion" unemployment in the sense captured by LAMBDA.

The triple–difference pre–period slope coefficient ($\varphi = 0.003$, statistically indistinguishable from zero) addresses one potential concern with this interpretation: that above– and below–median states had structurally different elasticities to *state–level* unemployment before 2008. What φ establishes is that the two groups were not differentially sensitive to their own labor market conditions before the recession: the two groups were, at least on state–level unemployment sensitivity, comparable entering the recession. The pattern suggests that safety net generosity shapes fertility advantages through the recovery period, rather than through acute crisis–insulation, and the absence of pre–recession differential sensitivity means this divergence cannot be explained away as a structural feature of the two groups that predates the recession itself.

7.2. What The Safety Net Is Doing Instead: Uncertainty and Institutional Confidence

So if the safety net is not cushioning unemployment to stabilize fertility, what is it doing? I posit that we should now turn to those who retain employment; for these people, the issue is neither income loss, nor time needed to search for work, but rather uncertainty about future conditions. Safety nets may matter less because they replace lost income, and more because they alter how people perceive the consequences of a potential future shock to an individual's income or time. Individuals living in more generous states may feel that unemployment, illness, or financial instability would be less catastrophic if they were to occur, reducing the perceived long–term risks associated with parenthood, even among households not receiving transfers. Paired with the faith the children are valued by the state, as reflected by spending priorities, this would generate the pattern observed in the data: a strong positive relationship between baseline generosity and fertility, paired with weak evidence for a contemporaneous unemployment–cushioning interaction.

7.3. The Young People Driving the Gap

The event–study results further build on this argument of prolonged support. The fertility differential does not emerge sharply during, or even directly after, the 2008–2010 collapse. Instead, it widens gradually during the recovery years, stabilizes around 2012–2015, and narrows only towards the end of the panel. Compared to a crisis–cushioning story, which would predict the largest differences during the unemployment peaks, followed by convergence as labor markets stabilized, this pattern is almost the reverse. Currie and Schwandt (2014) provide a useful potential explanation, finding that unemployment shocks have their largest fertility effects on women aged 20–24, with persisting effects after the unemployment–rate period. That age cohort during the 2008–2010 spike would have entered prime first–birth years during the 2012–2015 period, right where the event study peaks. The delayed timing may point towards mechanisms affecting how young people go on to perceive institutional support for families and thus the feasibility of parenthood, differentiated by the support system they experienced during the formative unemployment spike. So, in 2014, when a 28 year old in a low–support state and a 28 year old in a high–support state faced the same improving national economy, the 28 year old in a high–support state may move forward with family formation, believing that a future negative shock will not be borne entirely by them, whereas the 28 year–old in the low–support state may have chosen to delay parenthood until they felt personally secure enough to weather potential future instability without much support.

This interpretation connects with Lesthaeghe's (1983) description of the Second Demographic Transition, which includes a depiction of a long–run shift in advanced economies toward family formation as a deliberate and contingent (on education, career establishment, financial stability, etc.) life choice, rather than a default. This reflects the view of children as a highly personal, and increasingly costly, life choice—the investment required to raise a child to a given quality standard has risen with increases in areas like education and childcare. The safety net result sits interestingly against the SDT framework rather than simply within it. If the SDT decline is driven by individualization — the retreat from collective obligation toward purely personal calculation — then the finding that generous states maintained higher fertility is not anomalous but expected: a more generous safety net is, by definition, a more collectivist institutional arrangement, one that partially re–socializes the costs and risks of parenthood. States that have sustained stronger collective commitments to family support may have partially arrested the individualization dynamic that the SDT predicts drives fertility down. The result does not contradict the SDT; it suggests that institutional design can moderate its fertility consequences. The connection to safety net generosity is not necessarily that individuals decide to have children based on marginal transfer amounts, but that living within a more generous institutional environment signals that some of the financial burden will be collectively absorbed, making the preconditions for parenthood feel achievable at a lower subjective or objective personal financial status.

7.4. The Kearney et al. Puzzle

Finally, I want to place this thesis's findings in conversation with Kearney, Levine, and Pardue's (2022) 'puzzle' regarding the post–2008 US fertility decline. A precise clarification is warranted: my DiD is identified from cross–state variation in pre–recession safety net generosity, not from a within–state shift that could explain the national time series. It therefore cannot resolve the Kearney et al. puzzle directly. What it can do is point toward a class of explanations — political, psychological, institutional — that their framework explicitly leaves open. Kearney et al. show that standard economic variables account for less than 10% of the sustained decline, and tentatively attribute the remainder to a broad cultural shift they cannot identify with their data. The state–level differences captured here suggest that perceived institutional support is a plausible candidate for part of that residual: the DiD differential may reflect not only the direct effect of cash transfers, but the cumulative signal sent by a state's policy environment about whether parenthood is a burden borne alone or a risk shared collectively. If that interpretation is correct, policies aimed at reversing the fertility decline by targeting access to IVF or reducing the cost of any single program may be addressing a narrower channel than the data suggest. That is much harder to legislate in a quick–fix bill — but it may provide a way forward, a question this thesis opens rather than closes.

7.5. Limitations

Several limitations should be kept in mind when interpreting the results.

The parallel trends test has limited statistical power. Because there are only four pre–period years, the F–test cannot rule out subtle or non–linear differential trends, and failure to reject is not proof. Also, the F–test is not directly applicable to the continuous DiD, and no analogous formal test is conducted for DELTA.

The safety net splice introduces a measurement discontinuity at 2008, where the series shifts from unadjusted nominal cashfood to COLA–adjusted real cashfood_rs. R12, which uses unadjusted nominal cashfood throughout, produces a significant positive LAMBDA (0.078**) — the opposite sign from a cushioning mechanism. This should not be read as evidence that there is an amplification mechanism which is real and merely obscured by the splice. The unadjusted series ignores purchasing power: a high nominal benefit in a high–cost state is worth less than the same number in a low–cost state, meaning R12's interaction at least partly captures

cost–of–living variation rather than safety net variation. The null LAMBDA holds across all eleven other specifications. This is the single most important threat to the LAMBDA finding: if the splice itself introduces systematic measurement error correlated with the unemployment–fertility interaction, the null could be an artifact. The strongest evidence against this interpretation is R11, which restricts the sample to 2008–2019 and uses cashfood_rs exclusively, thus eliminating the splice entirely, and still produces a null LAMBDA (−0.002). If the splice were driving the null, removing it should change the estimate materially; it does not.
A related concern is that the preferred model's within–R^2 rises notably when restricted to the post–2007 period (R11: 0.681) or when year fixed effects are used (A6: 0.687), compared to the preferred specification (0.552–0.564). This suggests a meaningful share of the panel regression's identifying variation is concentrated in the recession–onset window rather than the recovery years where the institutional confidence mechanism is argued to operate. The null on LAMBDA is therefore most precisely estimated in the period least relevant to the proposed mechanism, which future work with longer post–recession panels could address.

The R13 result (0.019***) is interpreted in Section 7.1 as evidence that the safety net's fertility advantage operates through the recovery period rather than acute crisis buffering. An alternative reading is that high–generosity states (by and large coastal, urban, and high–FLFP) are structurally more sensitive to labor market conditions on fertility, independent of the safety net's causal role. State fixed effects absorb level differences, but not differences in the elasticity of fertility to economic conditions. Distinguishing these interpretations would require individual–level data on fertility intentions and economic expectations.

Safety net generosity may also proxy for broader cultural and institutional differences (i.e. public attitudes toward family formation, quality of public services, political culture) that collectively shape whether parenthood feels viable. This is less a confound than a clarification: the coefficients should be interpreted as the effect of a broader institutional environment rather than any specific transfer program.

Since state residency is not fixed, post–recession in–migration toward generous states could bias DELTA upward, though the gradual 2010–2015 build in the event study is more consistent with a structural effect than a compositional shift. The event study nonetheless cannot account for differential economic recovery timelines, and the mechanisms proposed above should accordingly be read as theoretically informed interpretations consistent with the observed patterns, rather than identified causal channels.

7.6. Directions for Future Research

An age–disaggregated event study tracing the cohort–specific timing of the fertility gap would directly test whether women aged 20–24 during the 2008–2010 unemployment spike — the group most sensitive to unemployment shocks (Currie & Schwandt, 2014) — account for a disproportionate share of the observed differential between high– and low–generosity states. This research design would better identify the institutional confidence mechanism by linking the aggregate pattern to the specific cohort whose formative economic experience the theory predicts should matter most.

The problem with the splice and data availability is not insignificant. The COVID–19 fertility shock, which generated dramatic and heterogeneous state–level responses in both fertility and social policy, may offer a second natural experiment for testing the same mechanism on a more recent and better–documented institutional landscape. With COVID–19, research could also include data on rent prices (which are more available in more recent years).

Future research could also combine individual–level fertility intention data with survey measures of institutional trust and perceived economic security to examine whether these beliefs predict delayed or foregone childbearing independently of objective economic conditions. Using individual–level survey data would allow a more direct test of the confidence channel proposed here and would help distinguish between the financial and psychological dimensions of how safety net generosity shapes family formation decisions. Repeated cross–sections from the National Survey of Family Growth, which offers data on child "wantedness," linked to state policy environments, offer one feasible path toward this kind of analysis.

The thesis also has policy implications that invite inspection. If the fertility advantage of generous states reflects the cumulative institutional signal of a state's policy environment rather than the marginal dollar value of any specific transfer, then incremental expansions of individual programs may be less effective at supporting family formation than broader, more visible commitments to institutional support for parents and children. Research comparing the fertility effects of targeted, means–tested expansions against universal or near–universal programs (i.e. paid parental leave, subsidized childcare, housing support) would help distinguish between these mechanisms and provide more actionable guidance for policymakers concerned with the fertility decline.

Finally, there is the cross–national dimension of this question. The institutional confidence mechanism proposed here maps naturally onto comparative welfare state research: if the cumulative policy signal shapes whether parenthood feels viable, Scandinavian countries — which historically combined high fertility relative to their income levels with generous, universal social provision — offer an interesting test case. A comparative design examining whether cross–national variation in welfare state generosity predicts fertility differentials like those documented here at the state level would both provide support the mechanism and clarify whether the US findings reflect something specific to the American context of thin and fragmented social provision, or a more general relationship between institutional support and family formation.

8. Conclusion

This thesis investigated whether state safety net generosity shaped the fertility response to the Great Recession, with a secondary question of whether that relationship operated through cushioning the immediate effect of unemployment on fertility. Using a 51–state panel spanning 2003–2019 and four complementary empirical strategies (a binary DiD with an event–study extension, a continuous DiD, a fixed–effects panel regression estimating LAMBDA, and a triple–difference specification) the results clearly address the first question: states that entered the Great Recession with more generous safety nets experienced significantly stronger post–2008 fertility trajectories. The continuous DiD and static DiD estimates remain positive, stable, and statistically significant across specification, while the event study shows a persistent fertility differential that builds gradually during the recovery years before decreasing slightly. On the other hand, the evidence for the unemployment cushioning mechanism is much weaker, and at best suggestive rather than conclusive. LAMBDA is null across eleven out of twelve specifications, and the triple–difference estimates provide only limited suggestive support for the mechanism.

Importantly, these two sets of estimates are not identifying the same type of variation. The DiD estimates are identified from between–state differences in pre–recession safety net generosity, which reflects decades of political choices and is determined prior to the recession. The LAMBDA estimates, on the other hand, rely on within–state year–to–year changes in unemployment and safety net generosity, which occur alongside the recession itself. This identifying variation is both smaller in magnitude and more vulnerable to endogeneity, since states often expand support programs when unemployment rises. The null on LAMBDA therefore reflects the difficulty of isolating an unemployment–cushioning channel through within–state panel variation during a major, complex macroeconomic shock. While I treat the null as a true null — particularly given the weakness of the triple–differences result, which is less vulnerable to these endogeneity concerns — it remains possible that the null is a methodological artifact and that the unemployment–cushioning mechanism is a key channel.

The magnitude of the DiD estimates is economically meaningful: a coefficient of approximately 0.035 log points translates to roughly 2.2 births per 1,000 women at the sample GFR, and the estimate changes very little when demanding controls are added, strengthening rather than attenuating the result. As Acemoglu and Finkelstein (2008) caution, however, state generosity likely captures a broader institutional environment that cannot be fully separated from the observable financial generosity. More generous states may differ not only in social welfare spending but also in political culture, public services, and social attitudes. The estimates should therefore be interpreted less as the isolated causal effect of a particular transfer program and more as the effect of living within a broader environment that produces stronger social support systems, as developed in Section 7.

The primary contribution of this thesis is to offer perceived institutional support as a candidate explanation for part of the post–2008 fertility decline that standard economic variables cannot account for (Kearney, Levine & Pardue, 2022). Several limitations apply, as discussed above, and the mechanisms proposed remain theoretically informed interpretations rather than identified causal channels. Future research could build on these findings through an age–disaggregated event study tracing cohort–specific fertility timing by state safety net level, and by combining individual–level fertility intention data with survey measures of institutional trust to test the confidence channel more directly.

References

Acemoglu, Daron, and Amy Finkelstein. 2008. "Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector." *Journal of Political Economy* 116 (5): 837–880.

Adesete, Ahmed and Piano, Clara, Unexpected Inflation and Fertility Rate in the US (November 01, 2025). Available at SSRN: https://ssrn.com/abstract=6038934 or http://dx.doi.org/10.2139/ssrn.6038934

Almond, Douglas, Hilary W. Hoynes, and Diane Whitmore Schanzenbach. 2011. "Inside the War on Poverty: The Impact of Food Stamps on Birth Outcomes." *Review of Economics and Statistics* 93 (2): 387–403

Angrist, Joshua D., and Jörn–Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton: Princeton University Press.

Balingit, Moriah, and Associated Press. 2025. "DOGE Could Worsen America's Child Care Crisis with Cuts to Programs." *Fortune*, April 2, 2025. https://fortune.com/2025/04/02/trump–doge–cuts–head–start–child–care–crisis–parents–jobs–economy/.

Baughman, Reagan, and Stacy Dickert–Conlin. 2009. "The Earned Income Tax Credit and Fertility." *Journal of Population Economics* 22 (3): 537–563.

Beaujouan, Éva. 2023. "Delayed Fertility as a Driver of Fertility Decline?" In *The Demography of Transforming Families*, edited by Robert Schoen, 41–63. Springer.

Becker, Gary S. 1960. "An Economic Analysis of Fertility." In *Demographic and Economic Change in Developed Countries*, edited by Universities–National Bureau Committee for Economic Research, 209–240. Princeton: Princeton University Press.

Becker, Gary S., and H. Gregg Lewis. 1973. "On the Interaction Between the Quantity and Quality of Children." *Journal of Political Economy* 81 (2): S279–S288.

Black, Dan A., Natalia Kolesnikova, Seth G. Sanders, and Lowell J. Taylor. 2013. "Are Children 'Normal'?" *Review of Economics and Statistics* 95 (1): 21–33.

Board of Trustees, Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds. 2025. *2025 Annual Report of the Boards of Trustees of the Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds*. https://www.cms.gov/oact/tr/2025

Brooks, Abigail, and Lovelace, Berkeley, Jr. 2025. "CDC's IVF Team Gutted Even as Trump Calls Himself the 'Fertilization President.'" *NBC News*. https://www.nbcnews.com/health/health–news/cdcs–ivf–team–gutted–even–trump–calls–fertilization–president–rcna199261

Brown, Sarah S., and Leon Eisenberg, eds. 1995. *The Best Intentions: Unintended Pregnancy and the Well–Being of Children and Families*. Washington, DC: National Academy Press.

Bureau of Labor Statistics. 2012. *The Recession of 2007–2009*. BLS Spotlight on Statistics. U.S. Department of Labor. https://www.bls.gov/spotlight/2012/recession/pdf/recession_bls_spotlight.pdf

Bureau of Labor Statistics. 2018. "Great Recession, Great Recovery? Trends from the Current Population Survey." *Monthly Labor Review*. U.S. Department of Labor. https://www.bls.gov/opub/mlr/2018/article/great–recession–great–recovery.htm

Bureau of Labor Statistics. 2026. *Local Area Unemployment Statistics*. U.S. Department of Labor. https://www.bls.gov/lau/

Bureau of Labor Statistics. 2026. *Current Population Survey: Annual Social and Economic Supplement*. U.S. Department of Labor. https://www.bls.gov/cps/

Bureau of Labor Statistics. 2026. "Table 14. Employment Status of the Civilian Noninstitutional Population by Sex, Race, Hispanic or Latino Ethnicity, Marital Status, and Detailed Age." *Current Population Survey*. U.S. Department of Labor. https://www.bls.gov/cps/tables.htm

Butz, William P., and Michael P. Ward. 1979. "The Emergence of Countercyclical U.S. Fertility." *American Economic Review* 69 (3): 318–328.

Centers for Disease Control and Prevention. 2026. *CDC WONDER: Natality Information*. National Center for Health Statistics. https://wonder.cdc.gov/natality.html

Congressional Budget Office. 2015. *Estimated Impact of the American Recovery and Reinvestment Act on Employment and Economic Output*. https://www.cbo.gov/publication/49958

Couillard, B. K. (2026, March 6). Build, Baby, Build: How Housing Shapes Fertility. Economics Department, University of Toronto.

Currie, Janet, and Hannes Schwandt. 2014. "Short– and Long–Term Effects of Unemployment on Fertility." *Proceedings of the National Academy of Sciences* 111 (41): 14734–14739.

Dettling, Lisa J., and Melissa S. Kearney. 2014. "House Prices and Birth Rates: The Impact of the Real Estate Market on the Decision to Have a Baby." *Journal of Public Economics* 110: 82–100.

Eissa, Nada, and Jeffrey B. Liebman. 1996. "Labor Supply Response to the Earned Income Tax Credit." *Quarterly Journal of Economics* 111 (2): 605–637.

Federal Housing Finance Agency. 2026. *All–Transactions House Price Index (Not Seasonally Adjusted, State–level)*. Annual data. https://www.fhfa.gov/data/hpi/datasets?tab=annual–data

Federal Reserve History. 2013. "The Great Recession." https://www.federalreservehistory.org/essays/great–recession–of–200709

Food and Nutrition Service. n.d. "SNAP – Supplemental Nutrition Assistance Program." U.S. Department of Agriculture. https://www.fns.usda.gov/snap

Forbes. 2025. "Elon Musk." *Forbes* Profile. https://www.forbes.com/profile/elon–musk/

Galbraith, Virginia L., and Dorothy S. Thomas. 1941. "Birth Rates and the Interwar Business Cycles." *Journal of the American Statistical Association* 36 (216): 465–476.

Galor, Oded, and David N. Weil. 1996. "The Gender Gap, Fertility, and Growth." *American Economic Review* 86 (3): 374–387.

Gascon, Charles S., and Ana Hernández Kent. 2025. "The Economics of Child Care: A State–Level Analysis." *On the Economy* (blog). Federal Reserve Bank of St. Louis. https://www.stlouisfed.org/on–the–economy/2025/may/economics–child–care–state–level–analysis

Gnoth, Christoph, Dörte Godehardt, Eva Godehardt, Peter Frank–Herrmann, and Günter Freundl. 2003. "Time to Pregnancy: Results of the German Prospective Study and Impact on the Management of Infertility." *Human Reproduction* 18 (9): 1959–1966.

Guzzo, Karen Benjamin, and Sarah R. Hayford. 2023. "Evolving Fertility Goals and Behaviors in Current U.S. Childbearing Cohorts." *Population and Development Review* 49 (1): 7–42.

Hamilton, Brady E., Joyce A. Martin, and Michelle J. K. Osterman. 2025. "Births: Provisional Data for 2024." *Vital Statistics Rapid Release*, no. 38, 1–10. https://dx.doi.org/10.15620/cdc/174587

Hamilton, Brady E., Michelle J. K. Osterman, and Elizabeth C. W. Gregory. 2026. "Births: Provisional Data for 2025." *Vital Statistics Rapid Release*, no. 43, 1–7. https://dx.doi.org/10.15620/cdc/252434

Hoynes, Hilary W., and Diane Whitmore Schanzenbach. 2009. "Consumption Responses to In–Kind Transfers: Evidence from the Introduction of the Food Stamp Program." *American Economic Journal: Applied Economics* 1 (4): 109–139.

Internal Revenue Service. 2025a. "EITC Income Limits, Maximum Credit Amounts." https://www.irs.gov/credits–deductions/individuals/earned–income–tax–credit

Internal Revenue Service. 2025b. "Child Tax Credit." https://www.irs.gov/credits–deductions/individuals/child–tax–credit

Jones, Larry E., and Michèle Tertilt. 2008. "An Economic History of Fertility in the U.S.: 1826–1960." In *Frontiers of Family Economics*, vol. 1, edited by Peter Rupert. Emerald Press.

Kearney, Melissa S. 2004. "Is There an Effect of Incremental Welfare Benefits on Fertility Behavior? A Look at the Family Cap." *Journal of Human Resources* 39 (2): 295–325.

Kearney, Melissa S., Phillip B. Levine, and Luke Pardue. 2022. "The Puzzle of Falling US Birth Rates since the Great Recession." *Journal of Economic Perspectives* 36 (1): 151–176.

Kearney, Melissa S., and Riley Wilson. 2018. "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence from the Fracking Boom." *Review of Economics and Statistics* 100 (4): 678–690.

Kile, Meredith. 2025. "RFK Jr. Claims 'Parents Aren't Having Children.'" *People*. https://people.com/rfk–jr–claims–parents–arent–having–children–11831749

Lesthaeghe, Ron. 1983. "A Century of Demographic and Cultural Change in Western Europe: An Exploration of Underlying Dimensions." *Population and Development Review* 9 (3): 411–435.

Lesthaeghe, Ron, and Dirk J. van de Kaa. 1986. "Twee Demografische Transities?" In *Bevolking: Groei en Krimp*, edited by Dirk J. van de Kaa and Ron Lesthaeghe, 9–24. Deventer: Van Loghum Slaterus.

Lovenheim, Michael F., and Kevin J. Mumford. 2013. "Do Family Wealth Shocks Affect Fertility Choices? Evidence from the Housing Market." *Review of Economics and Statistics* 95 (2): 464–475.

Martin, Joyce A., Brady E. Hamilton, Paul D. Sutton, Stephanie J. Ventura, T. J. Mathews, and Michelle J. K. Osterman. 2010. "Births: Final Data for 2008." *National Vital Statistics Reports* 59 (1). National Center for Health Statistics. https://www.cdc.gov/nchs/data/nvsr/nvsr59/nvsr59_01.pdf

Mincer, Jacob. 1963. "Market Prices, Opportunity Costs, and Income Effects." In *Measurement in Economics: Studies in Mathematical Economics and Econometrics in Memory of Yehuda Grunfeld*, edited by Carl F. Christ et al., 67–82. Stanford: Stanford University Press.

Moffitt, Robert A. 1998. "The Effect of Welfare on Marriage and Fertility." In *Welfare, the Family, and Reproductive Behavior*, edited by Robert A. Moffitt, 50–97. Washington, DC: National Academy Press.

NBER. 2010. "Business Cycle Dating Committee Announcement." National Bureau of Economic Research. https://www.nber.org/news/business–cycle–dating–committee–announcement–september–20–2010

Office of Family Assistance. n.d. "TANF: Temporary Assistance for Needy Families." U.S. Department of Health and Human Services. https://www.acf.hhs.gov/ofa/programs/tanf

Olesiuk, Shayna M., and Kathy R. Kalser. 2009. "The Sand States: Anatomy of a Perfect Housing Market Storm." *FDIC Quarterly* 3 (1): 30–32. https://www.fdic.gov/analysis/quarterly–banking–profile/fdic–quarterly/2009–vol3–1/vol3–1–sand–states.pdf

Osterman, Michelle J. K., Brady E. Hamilton, Joyce A. Martin, Anne K. Driscoll, and Claudia P. Valenzuela. 2025. "Births: Final Data for 2023." *National Vital Statistics Reports* 74 (1). National Center for Health Statistics. https://www.cdc.gov/nchs/data/nvsr/nvsr74/nvsr74–1.pdf

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. 2023. "What's Trending in Difference–in–Differences? A Synthesis of the Recent Econometrics Literature." *Journal of Econometrics* 235 (2): 2218–2244.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rodgers, Jonathan Schroeder, and Kari C.W. Williams. *IPUMS USA: Version 16.0* [dataset]. Minneapolis, MN: IPUMS, 2025. https://doi.org/10.18128/D010.V16.0.

Schaller, Jessamyn. 2016. "Booms, Busts, and Fertility: Testing the Becker Model Using Gender–Specific Labor Demand." *Journal of Human Resources* 51 (1): 1–29.

Schmidt, Lucie, Lara Shore–Sheppard, and Tara Watson. 2024. *State Safety Net Database*, version 2.1. Dataset. Brookings Institution, Washington, D.C.

Schneider, Daniel. 2015. "The Great Recession, Fertility, and Uncertainty: Evidence from the United States." *Journal of Marriage and Family* 77 (5): 1144–1156.

Schultz, T. Paul. 1973. "A Preliminary Survey of Economic Analyses of Fertility." *American Economic Review* 63 (2): 71–78.

Semega, Jessica L., Kayla R. Fontenot, and Melissa A. Kollar. 2017. *Income and Poverty in the United States: 2016*. Current Population Reports, P60–259. U.S. Census Bureau. https://www.census.gov/library/publications/2017/demo/p60–259.html

Shead, Sam. 2021. "Elon Musk Says 'Civilization Is Going to Crumble' If People Don't Have More Children." *CNBC*, December 7, 2021. https://www.cnbc.com/2021/12/07/elon–musk–civilization–will–crumble–if–we–dont–have–more–children.html

Simon, Curtis J., and Robert Tamura. 2009. "Do Higher Rents Discourage Fertility? Evidence from U.S. Cities, 1940–2000." *Regional Science and Urban Economics* 39 (1): 33–42.

Sobotka, Tomáš, Vegard Skirbekk, and Dimiter Philipov. 2011. "Economic Recession and Fertility in the Developed World." *Population and Development Review* 37 (2): 267–306.

Tax Policy Center. 2025. "Earned Income Tax Credit Parameters." https://www.taxpolicycenter.org

U.S. Bureau of Labor Statistics, Unemployment Rate [UNRATE], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/UNRATE, May 10, 2026.

U.S. Census Bureau. 2026a. *Real Median Household Income in the United States* [MEHOINUSA672N]. Retrieved from FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/MEHOINUSA672N

United States Census Bureau. 2026. *Small Area Income and Poverty Estimates (SAIPE) Program Datasets*. https://www.census.gov/programs–surveys/saipe/data/datasets.html

United States Census Bureau. *American Community Survey 1–Year Estimates, Table S1501: Educational Attainment*, 2005–2019 [dataset]. Washington, DC: U.S. Census Bureau, 2019. https://data.census.gov. (accessed via IPUMS)

United States Department of Labor. n.d. "Family and Medical Leave Act." https://www.dol.gov/agencies/whd/fmla

Willis, Robert J. 1973. "A New Approach to the Economic Theory of Fertility Behavior." *Journal of Political Economy* 81 (2): S14–S64.

Wolfers, Justin. 2006. "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results." *American Economic Review* 96 (5): 1802–1820.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2nd ed. Cambridge, MA: MIT Press.

World Bank, Fertility Rate, Total for the United States [SPDYNTFRTINUSA], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/SPDYNTFRTINUSA, May 10, 2026.

Yule, G. Udny. 1906. "On the Changes in the Marriage– and Birth–Rates in England and Wales During the Past Half Century; with an Inquiry as to Their Probable Causes." *Journal of the Royal Statistical Society* 69 (1): 88–132.

Appendix

Table A. Safety Net Generosity: Cashfood Values and Rankings, 2007–2009

State	2007		2008			Rank Shift 2007→2008	2009	
	cashfood	Rank	cashfood	cashfood_rs	Rank		cashfood	cashfood_rs
	(unadj.)	(unadj.)	(unadj.)	(COLA–adj.)	(unadj.)	(unadj.)	(unadj.)	(COLA–adj.)
Alabama	4.606	3	4.647	5.177	3	0	5.393	6.081
Alaska	8.261	51	8.213	7.732	50	1	9.187	8.563
Arizona	5.070	17	5.089	4.951	18	1	5.676	5.478
Arkansas	4.613	4	4.649	5.224	4	0	5.395	6.168
California	6.834	48	6.684	6.079	47	1	7.344	6.725
Colorado	5.047	16	5.067	4.920	16	0	6.045	5.926
Connecticut	5.302	24	5.311	4.863	24	0	6.094	5.584
Delaware	5.326	26	5.341	5.299	26	0	6.138	6.022
District of Columbia	6.045	43	6.171	5.516	44	1	7.018	6.252
Florida	5.187	22	5.205	5.114	21	1	5.953	5.901
Georgia	4.813	8	4.842	5.109	6	2	5.589	6.055
Hawaii	8.109	50	8.474	7.723	51	1	9.390	8.369
Idaho	4.998	15	5.014	5.270	15	0	5.762	6.039
Illinois	5.082	19	5.181	5.164	20	1	5.941	5.834
Indiana	4.912	14	4.931	5.328	12	2	5.721	6.153
Iowa	5.484	32	5.559	6.301	32	0	6.325	7.044
Kansas	5.578	35	5.605	6.177	34	1	6.395	6.929
Kentucky	4.789	5	4.819	5.347	5	0	5.566	6.255
Louisiana	4.868	9	4.960	5.343	14	5	5.716	6.210
Maine	5.715	38	5.684	5.810	36	2	6.426	6.599
Maryland	5.847	42	5.936	5.473	42	0	6.783	6.177
Massachusetts	6.268	45	6.224	5.888	45	0	7.014	6.703
Michigan	5.707	37	5.870	6.144	41	4	6.857	7.162

Minnesota	7.123	49	7.116	7.386	49	0	8.293	8.411
Mississippi	4.564	2	4.605	5.165	2	0	5.351	6.162
Missouri	4.911	13	4.939	5.527	13	0	5.686	6.248
Montana	5.640	36	5.627	5.872	35	1	6.475	6.845
Nebraska	5.314	25	5.315	5.851	25	0	6.111	6.666
Nevada	5.766	41	5.757	5.669	40	1	6.582	6.352
New Hampshire	5.720	39	5.689	5.372	38	1	6.620	6.391
New Jersey	5.442	28	5.491	5.138	28	0	6.345	5.720
New Mexico	5.482	31	5.688	5.992	37	6	6.463	6.725
New York	6.499	46	6.465	5.821	46	0	7.388	6.677
North Carolina	4.790	6	4.881	5.250	7	1	5.667	6.175
North Dakota	5.547	34	5.543	6.383	30	4	6.292	7.039
Ohio	5.742	40	5.723	6.187	39	1	6.560	7.020
Oklahoma	5.136	20	5.151	5.669	19	1	5.911	6.576
Oregon	5.365	27	5.388	5.520	27	0	6.151	6.158
Pennsylvania	5.076	18	5.087	5.135	17	1	5.835	5.876
Rhode Island	6.135	44	6.092	5.970	43	1	6.859	6.681
South Carolina	4.808	7	4.896	5.282	8	1	5.660	6.168
South Dakota	5.473	30	5.551	6.456	31	1	6.300	7.223
Tennessee	4.883	11	4.929	5.355	11	0	5.677	6.249
Texas	4.882	10	4.925	5.051	10	0	5.683	5.840
Utah	5.445	29	5.494	5.641	29	0	6.243	6.243
Vermont	6.609	47	6.698	6.630	48	1	7.530	7.472
Virginia	5.186	21	5.208	5.043	22	1	6.003	5.852
Washington	5.520	33	5.561	5.348	33	0	6.311	6.040
West Virginia	4.898	12	4.924	5.586	9	3	5.671	6.523
Wisconsin	4.443	1	4.494	4.815	1	0	5.291	5.623
Wyoming	5.288	23	5.298	5.520	23	0	6.540	6.748

Notes: cashfood is the unadjusted nominal combined cash–and–food benefit index (thousands of nominal dollars). cashfood_rs is the COLA–adjusted real version, available only from 2008 onward. Rank is based on unadjusted cashfood within each year (1 = least generous, 51 = most generous). Rank shift is the absolute change in unadjusted rank between 2007 and 2008; mean = 0.94, maximum = 6. Rows shaded yellow indicate rank shifts ≥ 3. The 2009 spike reflects ARRA–driven safety net expansions. Source: Schmidt et al. (2024), Brookings Safety Net Index v2.1.

Table B. Binary DiD — Full Coefficient Output (ATT)

Dependent variable: log general fertility rate. Extends Table 2. All specifications include state fixed effects and standard errors clustered at the state level.

	(S1)	*(S2)*	*(S3)*	*(S4)*
	Basic	*+ Time controls*	*+ State trends*	*+ Full controls*
ATT: Treatment group × post-2008	**0.027***	**0.027***	**0.030*****	**0.034*****
	(0.0153)	*(0.0153)*	*(0.0097)*	*(0.0081)*
Post-2008 indicator	−0.068***	0.039***	0.037***	0.019***
	(0.0094)	*(0.0085)*	*(0.0055)*	*(0.0058)*
National unemployment rate		−0.005***	−0.005***	−0.004***
		(0.0007)	*(0.0007)*	*(0.0008)*
Linear time trend		−0.012***		
		(0.0010)		
Log FHFA HPI (lagged)				0.105***
				(0.0074)
Female BA share, % (lagged)				−0.007***
				(0.0020)
Log median household income (lagged)				−0.030
				(0.0281)
N	867	867	867	816
R^2 within	0.217	0.479	0.749	0.834
Time controls	No	Yes	Yes	Yes
State-specific linear trends	No	No	Yes	Yes
Full controls (HPI, education, income)	No	No	No	Yes
Pre-trend $F(4,40) = 0.394$ [$p = 0.812$]				

Notes: Standard errors in parentheses, clustered at the state level. S3 and S4 replace the common time trend with state-specific linear trends following Wolfers (2006). Full controls (S4) add log FHFA HPI, female BA share, and log real median household income, all lagged one year. Pre-trend $F(4,40) = 0.394$, $p = 0.812$.

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Source: Author's calculations. CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); BLS LAUS (Bureau of Labor Statistics, 2026); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026); FHFA House Price Index (Federal Housing Finance Agency, 2026); ACS Educational Attainment (Ruggles et al., 2025); SAIPE Median Household Income (U.S. Census Bureau, 2026).

Table C. Continuous DiD — Full Coefficient Output (DELTA)

Dependent variable: log general fertility rate. Extends Table 3. All specifications include state fixed effects and standard errors clustered at the state level. C2 is the preferred specification.

	(C1)	*(C2)*	*(C3)*
	Basic	*+ Time controls*	*+ Full controls*
DELTA: Safety net generosity × post-2008	**0.035*** **	**0.035*** **	**0.032*** **
	(0.0101)	*(0.0101)*	*(0.0089)*
Post-2008 indicator	−0.055***	0.052***	0.039***
	(0.0072)	*(0.0044)*	*(0.0055)*
National unemployment rate		Included	Included
Linear time trend		Included	Included
Log FHFA HPI (lagged)			Included
Female BA share, % (lagged)			Included
Log median household income (lagged)			Included
N	867	867	816
R^2 within	0.240	0.502	0.591
Time controls	No	Yes	Yes
Full controls (HPI, education, income)	No	No	Yes

Notes: Standard errors in parentheses, clustered at the state level. Treatment variable is each state's demeaned 2008 cashfood_rs value (COLA-adjusted, thousands of 2022 USD). DELTA > 0 indicates more generous states maintained better post-2008 fertility trajectories. The post-2008 indicator coefficient changes sign from C1 to C2 because the national unemployment rate and time trend absorb the common post-period level shift. Full controls add log FHFA HPI, female BA share, and log real median household income, all lagged one year.

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Source: Author's calculations. CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); BLS LAUS (Bureau of Labor Statistics, 2026); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026); FHFA House Price Index (Federal Housing Finance Agency, 2026); ACS Educational Attainment (Ruggles et al., 2025); SAIPE Median Household Income (U.S. Census Bureau, 2026).

Table D. Year Fixed Effects Coefficients — Specification A6

Dependent variable: log general fertility rate. A6 replaces the national unemployment rate and linear time trend with year fixed effects. Reference year = 2003. All other A6 coefficients reported in Table 4.

	Coefficient	*(SE)*
2004	0.000	*(.)*
2005	0.000	*(0.003)*
2006	0.029***	*(0.005)*
2007	0.034***	*(0.007)*
2008	0.021**	*(0.008)*
2009	0.008	*(0.005)*
2010	0.013	*(0.018)*
2011	0.003	*(0.022)*
2012	−0.006	*(0.018)*
2013	−0.021	*(0.016)*
2014	−0.022	*(0.015)*
2015	−0.038***	*(0.012)*
2016	−0.060***	*(0.014)*
2017	−0.093***	*(0.014)*
2018	−0.123***	*(0.015)*
2019	−0.149***	*(0.018)*
N	816	
R^2 within	0.687	

Notes: Standard errors in parentheses, clustered at the state level. Year FE absorb all common time variation, including the secular fertility decline. The accelerating negative trend from 2015 onward (−0.038*** to −0.149***) reflects the post-2014 phase of the fertility decline documented by Kearney, Levine & Pardue (2022), which is largely unexplained by standard economic variables.

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Source: Author's calculations. CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); BLS LAUS (Bureau of Labor Statistics, 2026); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026).

Table E. Robustness Checks I — Specification Alternatives (R1–R5)

Dependent variable: log general fertility rate. All specifications include state fixed effects and standard errors clustered at the state level. Primary estimand is LAMBDA (unemployment × safety net interaction). Dashes (—) indicate coefficient not included in that specification.

	(R1)	(R2)	(R3)	(R4)	(R5)
	Contemp.	Cash-only (2009–19)	HPI as shock	Drop DC & AK	Drop 2008–09
Primary estimand					
LAMBDA: Unemployment × safety net	**0.016**	**0.023**		**0.068**	**0.020**
	(0.068)	(0.036)		(0.065)	(0.075)
LAMBDA: HPI × safety net			−0.159		
			(0.129)		
Unemployment					
Log state unemployment rate	−0.069***			−0.049***	−0.041*
	(0.022)			(0.014)	(0.022)
Log safety net generosity (demeaned)	−0.082	0.008	−0.058	−0.008	−0.025
	(0.057)	(0.051)	(0.042)	(0.051)	(0.061)
Log unemp × safety net (no demean)		0.023			
		(0.036)			
Log unemp (lagged, demeaned)		0.011			
		(0.016)			
Housing					
Log FHFA HPI (lagged, demeaned)			0.084***		
			(0.029)		
Cash-only safety net (R2)					
Log cashfood_rs (lagged)		0.008			
		(0.051)			
Time controls					
National unemployment rate	0.013***	−0.030***	0.002*	0.006***	0.006*
	(0.004)	(0.004)	(0.001)	(0.001)	(0.004)
Linear time trend	−0.007***	−0.030***	−0.009***	−0.008***	−0.008***
	(0.001)	(0.003)	(0.001)	(0.001)	(0.001)
N	867	561	816	784	714
R^2 within	0.473	0.683	0.557	0.552	0.516

Notes: R1: contemporaneous (no lag) unemployment and safety net, N = 867. R2: cash-only safety net (cashfood_rs), restricted to 2009–2019 due to data availability, N = 561. R3: replaces unemployment with FHFA HPI as the shock variable. R4: drops DC and Alaska. R5: drops acute recession years 2008–2009. Standard errors clustered at the state level throughout.
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Source: Author's calculations. CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); BLS LAUS (Bureau of Labor Statistics, 2026); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026); FHFA House Price Index (Federal Housing Finance Agency, 2026).

Table F. Robustness Checks II — Controls and Leverage (R7–R13)

Dependent variable: log general fertility rate. All specifications include state fixed effects and standard errors clustered at the state level. R12 replaces COLA-adjusted cashfood_rs with unadjusted nominal cashfood throughout (splice robustness check). R13 replaces state unemployment × safety net with national unemployment × safety net.

	(R7)	(R8)	(R9)	(R10)	(R11)	(R12)	(R13)
	Kitchen sink	Drop CA, TX, NY	2-year lag	FLFP control	Post-2007 only	Unadj. cashfood	Natl. unemp × SN
Primary estimand							
LAMBDA: Unemp × safety net (demeaned)	**0.059**	**0.066**	**0.044**	**0.059**	**−0.002**	—	—
	(0.062)	(0.064)	(0.067)	(0.065)	(0.055)		
LAMBDA: Unemp × unadj. cashfood (R12)	—	—	—	—	—	0.078**	—
						(0.032)	
LAMBDA: Natl. unemp × safety net (R13)	—	—	—	—	—	—	0.019***
							(0.006)
Unemployment & safety net							
Log state unemp rate (lagged, demeaned)	−0.030**	−0.049***	—	−0.050***	−0.010	−0.038**	−0.055***
	(0.015)	(0.015)		(0.014)	(0.011)	(0.018)	(0.013)
Log safety net generosity (lagged, demeaned)	−0.015	−0.039	—	−0.011	−0.011	—	−0.135*
	(0.046)	(0.044)		(0.048)	(0.038)		(0.069)
Log unadj. cashfood (lagged, demeaned, R12)	—	—	—	—	—	−0.104	—
						(0.068)	
Log unemp (2-year lag, demeaned, R9)	—	—	−0.026**	—	—	—	—
			(0.012)				
Log cashfood (2-year lag, demeaned, R9)	—	—	0.040	—	—	—	—
			(0.050)				
Additional controls							
Log FHFA HPI (lagged)	0.055*	—	—	—	—	—	—
	(0.031)						
Female LFPR, % (lagged)	0.002	—	—	0.001	—	—	—
	(0.002)			(0.002)			
Log median household income (lagged)	0.013	—	—	—	—	—	—
	(0.039)						
Female BA share, % (lagged)	−0.004	—	—	—	—	—	—
	(0.003)						
Time controls							
National unemployment rate	0.004***	0.007***	0.000	0.006***	−0.005***	0.006***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Linear time trend	−0.006***	−0.008***	−0.010***	−0.008***	−0.014***	−0.007***	−0.008***
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
N	816	768	765	816	612	816	816
R^2 within	0.565	0.542	0.600	0.553	0.681	0.559	0.557

Notes: R7: kitchen sink with full control battery (HPI, FLFP, income, education). R8: drops California, Texas, and New York. R9: two-year lag on all variables. R10: FLFP control only. R11: post-2007 subsample (2008–2019), cashfood_rs throughout. R12: unadjusted nominal cashfood used throughout the full panel as a splice robustness check; positive significant LAMBDA (0.078**) reflects sensitivity to the measurement discontinuity at 2008. R13: national unemployment × safety net interaction; positive coefficient (0.019***) indicates generous states see larger relative fertility declines in high-national-unemployment years, consistent with the recovery-period mechanism rather than acute crisis buffering.

* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Source: Author's calculations. CDC WONDER Natality Database (Centers for Disease Control and Prevention, 2026); Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024); BLS LAUS (Bureau of Labor Statistics, 2026); FRED UNRATE (Federal Reserve Bank of Saint Louis, 2026); FHFA House Price Index (Federal Housing Finance Agency, 2026); BLS CPS Table 14 (Bureau of Labor Statistics, 2026); ACS Educational Attainment (Ruggles et al., 2025); SAIPE Median Household Income (U.S. Census Bureau, 2026).

Table G: State Split for Control vs Treatment Group

statename	sn_pre~g	sn_group
Wisconsin	4.443199	Bottom 50% — control
Mississippi	4.564326	Bottom 50% — control
Alabama	4.606362	Bottom 50% — control
Arkansas	4.612889	Bottom 50% — control
Kentucky	4.78861	Bottom 50% — control
North Carolina	4.789591	Bottom 50% — control
South Carolina	4.808455	Bottom 50% — control
Georgia	4.813187	Bottom 50% — control
Louisiana	4.868	Bottom 50% — control
Texas	4.881966	Bottom 50% — control
Tennessee	4.883226	Bottom 50% — control
West Virginia	4.897608	Bottom 50% — control
Missouri	4.911473	Bottom 50% — control
Indiana	4.911906	Bottom 50% — control
Idaho	4.997639	Bottom 50% — control
Colorado	5.046542	Bottom 50% — control
Arizona	5.069544	Bottom 50% — control
Pennsylvania	5.075728	Bottom 50% — control
Illinois	5.082325	Bottom 50% — control
Oklahoma	5.135867	Bottom 50% — control
Virginia	5.186447	Bottom 50% — control
Florida	5.187022	Bottom 50% — control
Wyoming	5.288332	Bottom 50% — control
Connecticut	5.302163	Bottom 50% — control
Nebraska	5.313586	Bottom 50% — control
Delaware	5.326058	Top 50% — treatment
Oregon	5.364726	Top 50% — treatment
New Jersey	5.441905	Top 50% — treatment
Utah	5.445173	Top 50% — treatment
South Dakota	5.472831	Top 50% — treatment
New Mexico	5.481645	Top 50% — treatment
Iowa	5.484232	Top 50% — treatment
Washington	5.52048	Top 50% — treatment
North Dakota	5.546838	Top 50% — treatment
Kansas	5.578455	Top 50% — treatment
Montana	5.639538	Top 50% — treatment
Michigan	5.706932	Top 50% — treatment
Maine	5.714535	Top 50% — treatment
New Hampshire	5.720069	Top 50% — treatment
Ohio	5.741938	Top 50% — treatment
Nevada	5.765722	Top 50% — treatment
Maryland	5.847023	Top 50% — treatment
District of Columbia	6.044981	Top 50% — treatment
Rhode Island	6.134583	Top 50% — treatment
Massachusetts	6.267898	Top 50% — treatment
New York	6.498745	Top 50% — treatment
Vermont	6.608941	Top 50% — treatment
California	6.834262	Top 50% — treatment
Minnesota	7.122951	Top 50% — treatment
Hawaii	8.109087	Top 50% — treatment
Alaska	8.261219	Top 50% — treatment

Source: Author's calculations. Brookings Safety Net Index v2.1 (Schmidt, Shore-Sheppard and Watson, 2024).

AI Appendix

In the process of writing this thesis, I used AI for three purposes: aiding with coding, polishing the text and visuals, and supporting the literature search process.

I used Claude primarily, as well as ChatGPT (ChatGPT only for triple–checking code). Both tools were used to clarify the meaning of coding errors and to assist in debugging throughout the empirical work. For the literature review, I used Claude to identify potentially relevant research areas and authors, especially by finding commonly–cited authors in the papers I had pre–selected. Still, all sources were independently verified and accessed through academic databases. For the written text, Claude was occasionally used to improve flow and clarity by suggesting changes to wording and sentence/paragraph structure.

On the coding side, AI tools proved genuinely valuable in accelerating and deepening the empirical work: because I could move faster, I was able to go further in my analysis than I otherwise would have been able. For the literature review, I remained cautious about over–relying on Claude to summarize texts given the known risk of hallucinated or incorrectly attributed sources. For text polish, AI assistance was helpful, though not without its limitations, as, when lacking context, Claude could delete things or rephrase thing in a way I did not appreciate.

To mitigate these risks, I cross–checked coding suggestions between models and always benchmarked results against the literature, verified all literature references independently, and reviewed all textual suggestions individually. Crucially, AI tools were not used in the production of original ideas — all arguments, interpretations, and conclusions are my own (correct or not!).

This has been the most I have ever engaged with AI, and it has been a great learning opportunity: AI tools offer real efficiency gains, particularly for technical work. However, they are not creative, cannot make “new” connections or ideas, and have no human judgment. I think that AI has incredible research applications, but it is also more than capable of creating very polished AI–slop, with no new contribution to knowledge. AI is a tool, but humans must use it like one and not underestimate the importance of original thought.