

NBER WORKING PAPER SERIES

IS THE IPHONE BIRTH CONTROL? CAUSAL EVIDENCE FROM AT&T'S 2007–2011
CARRIER MONOPOLY

Caitlin K. Myers
Ezekiel Hooper

Working Paper 35310
<http://www.nber.org/papers/w35310>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2026

This research project began when Ezekiel Hooper was an undergraduate at Middlebury College; he currently is employed by Accenture. We thank Daniel Dench and Phillip Levine for helpful comments. Claude Code (Anthropic) was used to assist with coding; the authors reviewed and are responsible for all code, analysis, and results. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Caitlin K. Myers and Ezekiel Hooper. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Is the iPhone Birth Control? Causal Evidence from AT&T's 2007–2011 Carrier Monopoly
Caitlin K. Myers and Ezekiel Hooper
NBER Working Paper No. 35310
June 2026
JEL No. J13, J18, O33

ABSTRACT

The U.S. general fertility rate has fallen by 22% since 2007, a sustained decline not readily explained by economic conditions, contraceptive use, housing or childcare costs, or other commonly cited factors. We assess the potential role of a different shock: the diffusion of the smartphone. The U.S. rollout of the iPhone, the first modern smartphone, provides a natural experiment: from June 2007 through February 2011, the device was sold only on AT&T, allowing us to identify its effect from variation in AT&T's mobile broadband coverage. Entropy-balanced Poisson and synthetic difference-in-differences event studies imply that access to the iPhone reduced births by 4.5–8.0% at ages 15–19 and 3.2–6.6% at ages 20–24, with statistically significant but smaller declines among older cohorts. Placebo analyses applied to Verizon and Sprint's pre-2011 coverage footprint are null. Taken together, these cohort effects imply that the diffusion of the iPhone deepened the decline in births among women under 30 while suppressing the rise in births among older women. Overall, the diffusion of the iPhone explains 33–52% of the decline in the general fertility rate among women aged 15–44. National-survey evidence on time use and sexual behavior is consistent with the iPhone reducing in-person interactions, increasing pornography use, and reducing sexual frequency.

Caitlin K. Myers
Middlebury College
Department of Economics
and NBER
cmyers@middlebury.edu

Ezekiel Hooper
Middlebury College
zekevt@gmail.com

1 Introduction

US births are plummeting. The general fertility rate (GFR) — annual live births per 1,000 women aged 15–44 — remained roughly constant at 65 to 70 from 1980 through 2007, then began a sharp decline that has continued for nearly two decades (Kearney et al., 2022). By 2024 the GFR had fallen to 54, a 22% decline over 17 years (Figure 1, Panel A). For a while, the conventional wisdom was that this simply reflected the effects of the Great Recession. US births are well-known to be procyclical (Kearney and Levine, 2015; Dettling and Kearney, 2023), the timing of the break aligned squarely with the onset of the recession, and contemporaneous observers expected fertility to rebound with the recovery (Sobotka et al., 2011; Cherlin et al., 2013; Congressional Budget Office, 2014). Nearly two decades on, that rebound has not come: the decline continued through the long economic expansion of 2010–2019, through the COVID dip and its aftermath, with declines visible across the same subgroups (Kearney and Levine, 2021; Kearney et al., 2022). A growing literature in economics, demography, and public health has tried — and so far largely failed — to identify a period-specific cause (Kearney and Levine, 2022; Kearney et al., 2022; Stone, 2018; Adashi et al., 2026).

Any candidate explanation must satisfy three demanding properties at once. It must be *big*: any factor that explains the break must account for a roughly 20% drop in the GFR over 15 years. It must be *dated*: any factor that explains the break must itself have shifted meaningfully around 2007 (Kearney and Levine, 2022). And it must be *broadly distributed*: Kearney et al. (2022) decompose the post-2007 fall and find that essentially every demographic subgroup contributes, by education, race/ethnicity, marital status, and parity, ruling out narrow group-specific stories. Within this broad pattern, the decline is sharpest in unintended births and at younger ages (Buckles et al., 2025). Between 2007 and 2024, U.S. birth

rates fell by 70% at ages 15–19 and 47% at ages 20–24, but only 7% at ages 30–34, while rates at 35–39 actually rose by 14% (Figure 1, Panel B). Kearney and Levine (2025) conclude that the decline reflects a cohort-level shift in priorities and name email, the iPhone, and social media as candidate “exogenous shocks that changed major facets of work and socializing across cohorts,” along with rising inequality and deindustrialization—but treat the empirical magnitude of each as an open question.

We focus on the role of the iPhone, which was launched in June 2007, diffused rapidly across every demographic group, and is consequential for how young adults spend their time. A descriptive literature in developmental psychology—most prominently Twenge (2017); Twenge and Park (2019); Twenge et al. (2019) and Haidt (2024)—documents that the cohorts whose adolescence overlapped with smartphone diffusion spend sharply less time on face-to-face socializing, dating, paid work, drinking, and driving, and that they report fewer sexual partners and rising sexual inactivity in young adulthood. Causal evidence supports the in-person-displacement channel: Allcott et al. (2020) document via RCT that four-week Facebook deactivation increases offline socializing with family and friends, and Braghieri et al. (2022) identify negative effects of Facebook access on student mental health from its staggered U.S. college rollout. This evidence extends to adults: Twenge et al. (2017) and Ueda et al. (2020) document declines in sexual frequency and rising sexual inactivity across the full 18–44 age range, not only among teens. Twenge (2023) predicts on the basis of these behavioral shifts that smartphone-shaped cohorts will produce fewer children than any previous generation in North American history.

While these correlations provide compelling support for an iPhone story, rigorous causal identification has been slow to arrive. Hudson and Moscoso Boedo (2026) document a synchronous global break in teen fertility around 2007 across

128 countries and attribute it to smartphone-driven displacement of in-person peer interaction. Their U.S. design treats the 2007 iPhone launch as a salient marker for the broader digital-era bundle rather than as the operative cause; we exploit its 2007–2011 AT&T-only U.S. launch to isolate the iPhone-specific channel and trace its incidence across age, race, parity, marital status, and education.

To empirically test the role of the iPhone, we leverage a different natural experiment arising from AT&T’s carrier monopoly. From its launch in June 2007 through February 2011, the iPhone was available in the United States only to subscribers of AT&T, so that a county’s exposure to the iPhone during these years was governed by the reach of AT&T’s mobile (cellular) broadband network—the infrastructure that carried smartphone data, as distinct from fixed wireline residential broadband. We compare counties with near-universal AT&T coverage to counties with little or none over 2003–2011, estimating age-specific birth-rate effects with two complementary designs: an entropy-balanced Poisson event study (Hainmueller, 2012) that reweights controls to match the treated covariate profile, and a synthetic difference-in-differences estimator (Arkhangelsky et al., 2021) that reweights controls and pre-treatment years to match the treated pre-iPhone outcome trajectory.

Both estimators imply large, statistically significant declines in births to young women. The post-gestation ATT ranges from -4.5 to -8.0% at ages 15–19 and -3.2 to -6.6% at ages 20–24 (the entropy-balanced Poisson at the lower-magnitude end, SDID at the higher), with smaller effects at older ages. Scaled to the U.S. county universe, these estimates imply the iPhone accounts for between 33 and 52% of the 2007–2011 decline in the general fertility rate. The pattern is similar across race, parity, marital status, and education, with the exception of Black women, for whom we estimate no effect. We additionally offer multiple placebo tests, including

placebo-in-time estimates that artificially assign the introduction of the iPhone to earlier years as well as tests using coverage by Sprint or Verizon—carriers that did not offer the iPhone until 2011—over the same window as that should register no effect.

Our findings supply the empirical magnitude Kearney and Levine (2025) declared an open question: the iPhone, and the smartphone era it inaugurated, materially accelerated the post-2007 U.S. fertility decline.

2 The AT&T 2007–2011 iPhone carrier monopoly

On January 9, 2007, Steve Jobs announced the iPhone at the Macworld convention in San Francisco (Apple Inc., 2007b). On the same day Apple named Cingular Wireless—then the largest U.S. wireless carrier, with roughly 58 million subscribers, in the process of being rebranded as AT&T following its parent’s 2006 acquisition of BellSouth—as Apple’s exclusive U.S. carrier partner for the device (Apple Inc., 2007a). Court filings later established that the underlying agreement granted AT&T five years of exclusivity over iPhone distribution in the United States (Patel, 2010).

The iPhone went on sale at 6:00 p.m. local time on Friday, June 29, 2007, through Apple’s 164 U.S. retail stores and AT&T retail locations, priced at \$499 for the 4 GB model and paired with AT&T service plans starting at \$59.99 per month on a mandatory two-year contract (Apple Inc., 2007c,d). Demand at launch was substantial, with multi-day queues at flagship stores widely reported in the contemporary press (CNN Money, 2007). For the four years that followed, the only legitimate means of using an iPhone in the United States was through an AT&T wireless contract.¹ The exclusivity was preserved through the iPhone 3G

¹While it was technologically possible to “jailbreak” an iPhone so that it could be used on another network, the iPhone’s radio was GSM-only, which made it compatible only with AT&T or T-Mobile networks. Contemporaneous estimates placed total jailbreak prevalence at 9% of U.S.

(2008), iPhone 3GS (2009), and iPhone 4 (2010), so that U.S. iPhone ownership was effectively coextensive with AT&T subscription throughout the window.

The smartphone landscape evolved during the AT&T-exclusive iPhone window. The first commercially available Android phone, the T-Mobile HTC Dream (marketed as the G1), launched on October 22, 2008 (T-Mobile USA, 2008). T-Mobile’s network, however, largely overlapped AT&T’s: in our sample, T-Mobile coverage averages 83.2% of the population in treated counties but only 3.7% in control counties (Table 1). The G1 launch therefore did not appreciably expand smartphone access in counties outside the AT&T footprint. Substantial Android exposure in our control counties arrived with two later launches on broader-coverage carriers: the Sprint HTC Hero on October 11, 2009 (Sprint Nextel Corporation, 2009) and the Verizon Motorola Droid on November 6, 2009 (Verizon Wireless, 2009), the latter backed by a roughly \$100 million “Droid Does” marketing campaign that contrasted Android’s capabilities against the iPhone. AT&T was the last of the four major U.S. carriers to add an Android handset, doing so in March 2010 with the Motorola Backflip (AT&T Inc., 2010b). The AT&T-versus-non-AT&T contrast in our sample is therefore partially contaminated by the rival-carrier Android channel from late 2009 onward. We address this concern through robustness checks described in Section 5: restricting the sample to counties with limited T-Mobile coverage (eliminating any G1-channel leakage) and truncating the analysis to a 2003–2009 window that pre-dates the Sprint and Verizon Android launches by gestation timing.

Apple’s exclusive arrangement with AT&T ended on February 10, 2011, when Apple released a Verizon-compatible variant of the iPhone 4 (Apple Inc., 2011).

iPhone and iPod Touch users by August 2009 (Chen, 2009), and only a fraction of those devices were carrier-unlocked rather than jailbroken for other purposes. Verizon and Sprint ran the incompatible CDMA standard and could not host any iPhone until the CDMA-compatible iPhone 4 in February 2011. As we discuss in the next paragraph, T-Mobile’s network largely overlapped with AT&T’s, leaving a jailbroken iPhone in a control county with no rival network to escape to. We treat this leakage as quantitatively negligible.

The release predated the nominal expiration of the original five-year contract by approximately a year; the contractual terms permitting the early termination were never publicly disclosed. The proximate engineering constraint on the timing was the absence of CDMA-compatible iPhone hardware: Verizon operated a CDMA2000 network, while the iPhone had been GSM-only, and the Verizon iPhone 4 launched in February 2011 was the first iPhone built with a dual-mode Qualcomm MDM6600 baseband (The CPU Shack Museum, 2011). The end of exclusivity also followed a period of regulatory and political scrutiny of handset-exclusivity practices, including a 2009 Federal Communications Commission inquiry into mobile wireless competition (Federal Communications Commission, 2009), hearings in the Senate (U.S. Senate Subcommittee on Antitrust, Competition Policy and Consumer Rights, 2009; U.S. Senate Committee on Commerce, Science, and Transportation, 2009), and multiple antitrust class-action suits (AppleInsider, 2010).

The implication for our identification strategy is that from June 29, 2007 through February 10, 2011, the functional use of an iPhone in the United States required both an AT&T wireless contract and access to AT&T’s mobile broadband network. Although the original 2007 iPhone ran on AT&T’s slower 2G EDGE network, the iPhone 3G (2008) and every subsequent model required AT&T’s 3G network, and the iPhone’s defining smartphone functions—web browsing, email, the App Store and its apps, and mobile maps—were either unavailable or practically unusable without a fast data connection. A consumer’s likelihood of owning a functional iPhone therefore depended on whether AT&T’s 3G network reached her county. This is the mobile (cellular) infrastructure on which smartphones operate—and on which the mobile peer-time displacement that drives the literature’s behavioral mechanism takes place—distinct from the fixed wireline residential broadband (cable, DSL, fiber) that carries home Wi-Fi to laptops and televisions. We exploit this feature of

the rollout in the empirical design developed in [Section 5](#).

3 Data

We assemble data on mobile broadband coverage by carrier (NTIA, 2010), births by five-year age groupings (NCHS, 2025), and populations by age, sex, and race (SEER, 2026) to construct a county-by-year panel dataset capturing iPhone availability and birth rates. Additional controls and complementary national survey data on proximate fertility behaviors and time use are described below.

3.1 Mobile Broadband by Carrier

In 2009, the National Telecommunications and Information Administration (NTIA) launched the State Broadband Initiative (SBI) to expand broadband access. As part of this effort, the NTIA issued grants to each of the 50 states, the District of Columbia, and U.S. territories to support the semiannual collection of data on the availability, speed, and location of broadband services, classified as either fixed or wireless. Wireless broadband is high-speed internet access delivered over radio signals as opposed to a wired last mile. It includes mobile (cellular) broadband provided over cellular networks and, in some contexts, fixed wireless service delivered to a fixed location via radio links.

Under SBI, wireless broadband carriers voluntarily provided states with information on the geographic coverage of their services. With support from the NTIA, states validated these submissions through field tests including mobile speed tests (Rachfal, 2021). The NTIA supported grantees' data collection and validation efforts, and in 2011 launched the National Broadband Map and began publishing wireless broadband coverage data at the census block level. The underlying wireless data are produced by overlaying validated carrier coverage shapes on census blocks

and computing the percentage of each block’s area covered by each provider. The National Broadband Map was subsequently updated semiannually through June 30, 2014, at which point the American Recovery and Reinvestment Act (ARRA) grant money supporting the project expired.²

For our research design, we focus on wireless broadband coverage and reported maximum speeds as of December 31, 2010 for the four national cellular carriers operating at that time: AT&T, Sprint, T-Mobile, and Verizon.³ During this period, the SBI did not attempt to distinguish fixed wireless versus mobile broadband. However, AT&T, T-Mobile, and Verizon did not offer consumer fixed wireless broadband services as part of their nationwide broadband portfolios, implying that their reported wireless coverage reflected mobile (cellular) broadband service. Sprint was the only national carrier with a fixed wireless broadband deployment during this period, operating a WiMAX network in partnership with Clearwire; however, this deployment was geographically limited (Kruger and Gilroy, 2009). As a result, wireless broadband coverage reported by national carriers in the SBI data provides a close approximation to mobile (cellular) broadband availability during the study period. We therefore use the term “mobile broadband” to refer to wireless broadband

²The Federal Communications Commission (FCC) subsequently assumed responsibility for the National Broadband Map, and continues to collect and issue information on fixed and mobile broadband coverage through its Form 477 program. Like the SBI program, the Form 477 program also relies on carriers’ self-reports of coverage. Unlike the SBI program, Form 477 data have historically lacked systematic independent validation. A substantial literature documents that Form 477 therefore overstates broadband availability, particularly in rural areas (see, e.g., Ford, 2021; Nasr, 2018). Our research design seeks to mitigate concerns about measurement error by relying on coverage data at a single point in time collected and validated under the SBI.

³Between June 2007 and June 2010, AT&T Mobility expanded its consolidated footprint via the acquisitions of Dobson Communications Corporation (2007), Easterbrooke Cellular Corporation, Windstream Wireless, and Edge Wireless (2008), Centennial Communications Corp. (2009), and certain wireless network properties divested by Verizon Wireless following its acquisition of Alltel (2010) (AT&T Inc., 2010a, 2011). The December 2010 SBI snapshot reflects this fully consolidated footprint. A county whose AT&T coverage arrived only via one of these acquisitions therefore had iPhone availability for less than the full AT&T-exclusivity window, biasing our post-gestation ATT estimates toward zero. The appendix includes a robustness check using a narrower 2003–2009 outcome window, which removes the 2010–2011 birth years over which this attenuation operates.

reported by the four national carriers.

We rely on the second release of data published by the SBI (NTIA, 2010), which provide a validated snapshot of mobile broadband coverage on national networks as of December 31, 2010. This snapshot falls between 2007—when the iPhone was launched through an exclusive carrier agreement with AT&T—and 2011—when Apple began offering the iPhone on Sprint and Verizon. To aggregate mobile broadband coverage from the census block level at which it is reported by the SBI to the county level at which births, our primary outcome, are observed, we first code a block as covered by a given carrier if the SBI reports any wireless broadband shape for that carrier at the block at speeds exceeding 200 Kbps—the FCC’s standard definition of broadband at the time the iPhone was released, and in practice the lowest reporting threshold in the SBI wireless file, so the cut amounts to a “did the carrier report any service” indicator.⁴ Block-level coverage is then averaged unweighted to the block group, and block-group shares are averaged to the county weighted by 2000 Census block-group population. [Figure 2](#) maps county-level variation in AT&T coverage in late 2010, illustrating that there is substantial variation, particularly across rural areas. For instance, while all of South Dakota was covered by AT&T, none of Montana was.

Much of our analysis rests on classifying counties as “control,” defined as <10% of the population residing in block groups with AT&T coverage and “treated,” defined as >90% of the population residing in block groups with AT&T coverage. We adopt this wide separation so that the treated and control populations have qualitatively different iPhone exposure. Intermediate counties with 10–90% cover-

⁴Every reported row in the SBI wireless data corresponds to an advertised maximum download speed of at least 200 Kbps—SBI’s lowest `MaxAdDown` code—so a binary “any coverage” indicator captures whether the carrier reports broadband service in the block. The 200 Kbps cut also captures AT&T’s mobile network as deployed at the snapshot date: AT&T reports no coverage shapes at ≥ 3 Mbps anywhere in the country in the December 2010 SBI snapshot, reflecting that the carrier’s 4G LTE rollout did not begin until late 2011.

age are excluded from analyses that rely on binary classifications.⁵ Following this assignment, 1,399 counties are classified as control and 914 as treated, while 794 counties are excluded. [Table 1](#) presents summary statistics by county assignment (control, treated, or excluded) during the baseline period prior to the introduction of the iPhone. Treated counties are systematically more urban than controls: the average control county has 28.2% of its population in urban areas, compared to 66.4% in the average treated county. We address this imbalance through the SDID estimator’s pre-treatment reweighting of controls and through a complementary entropy-balanced Poisson specification, both described in [Section 5](#).

3.2 Births

We measure age-specific births at the county level using restricted-use natality microdata from the National Center for Health Statistics (NCHS), in which each record is a single birth carrying the mother’s age, race and ethnicity, marital status, educational attainment, parity, and county of residence at the time of birth (NCHS, 2025); births are assigned to the mother’s county of residence, not county of occurrence. We aggregate to counts of births by county, year, and five-year age band of the mother for women aged 15–44 and divide by mid-year population estimates of women in the same county-by-age-by-race cells produced by the Surveillance, Epidemiology, and End Results Program (SEER, 2026) to construct age-specific birth rates per 1,000 women. Because gestation runs roughly nine months, the first birth cohort whose conceptions could be exposed to the iPhone is 2008; we therefore treat 2007 as a pre-treatment year and 2008 as the first treated cohort throughout the analysis. Our log-rate panels are constructed as balanced 2003–2011 panels by county and five-year age band; counties with a zero-birth cell in any year fall out of the age-

⁵Appendix A illustrates that the results are not sensitive to alternative thresholds of 5/95% or 20/80%, or to estimating models relying on a continuous measure of coverage.

specific log-rate panel because the log is undefined at zero, but enter the level-rate and Poisson specifications, which accommodate zero counts. Heterogeneity analyses below disaggregate these rates further by the mother’s race, ethnicity, marital status, parity, and educational attainment. Education-stratified rows are estimated on the 2003–2010 window only, because NCHS dropped the unrevised-certificate maternal-education item starting in 2011, leaving education unreported for births in states still using the 1989 birth certificate format (see the note to [Figure 5](#)).

3.3 Controls

[Table 1](#) summarizes additional county-level control variables. These include the racial and age composition of women aged 15–44 from SEER (2026); the share of the county’s population residing in decennial-Census urban areas, harmonized across years by IPUMS NHGIS (Schroeder et al., 2025); a four-region Census division indicator; the county-level 2008 Republican presidential vote share (MIT Election Data and Science Lab, 2018); annual unemployment from the Bureau of Labor Statistics (BLS, 2025); poverty rates and median household income from the Census Bureau’s Small Area Income and Poverty Estimates program (SAIPE, 2024); the year-over-year change in the Federal Housing Finance Agency’s county-level Home Price Index (all-transactions; FHFA, 2026); and state exposure to parental involvement laws (Myers and Ladd, 2020) and two-trip mandatory waiting periods for abortion (Myers, 2021).

3.4 National survey data for descriptive mechanisms

The remaining mechanism outcomes are drawn from national publicly-available data sources that lack the geographic identifiers needed to link to county AT&T coverage; we use them only to construct descriptive national trends and not as inputs to the

causal design.

Time with non-household members. We use the American Time Use Survey (ATUS) for survey years 2003–2024 (BLS, 2025) to construct national descriptive measures of time spent with friends present and time spent alone for respondents aged 15–44. The ATUS asks each respondent to recount the prior day’s activities in chronological episodes and, for each episode records who, if anyone, was physically with the respondent.⁶ Our two reported measures are minutes per day with friends present (excluding episodes during work or class) and minutes per day alone (episodes in which the respondent reports no other person physically present). The 2020 ATUS year is omitted from descriptive trends by age group because the Bureau of Labor Statistics does not produce the standard ATUS final weight for that survey year due to field-work disruption during the Covid pandemic.⁷⁸

Sexual behavior. We use the National Survey of Family Growth (NSFG), a nationally representative cross-sectional survey conducted by the CDC’s National Center for Health Statistics, for survey cycles 2002 through 2017–2019 (NCHS, 2024). We pool the public-use female respondent files across cycles, restrict to women aged 15–44, and apply each cycle’s final survey weight. Our three reported measures are the share of women who had sexual intercourse in the past month, the share who used any contraception at last sex, and an integrated indicator equal to one for women who had recent sex without contraception.

⁶The ATUS does not ask this question when the activity is sleep, personal care, or intimate activities.

⁷The measure of time with friends uses `tuwho_code` 54 (“Friends”) and the measure of time alone uses `tuwho_code` 18 or 19 (both labeled “Alone”).

⁸The ATUS reports CBSA identifiers as well as county identifiers for counties with populations greater than 100,000. In the appendix we explore the feasibility of using these limited geographic identifiers to estimate Poisson event studies of outcomes, but conclude that the limited sample size, combined with measurement error introduced by collapsing many counties into larger spatial units, renders the approach of limited usefulness.

Psychological distress. We use the National Health Interview Survey (NHIS) for survey years 2002–2024, accessed via IPUMS Health Surveys (Blewett et al., 2024), to construct the share of adults aged 18–44 with a Kessler-6 (K6) score of at least 13, the standard threshold for serious psychological distress. We reconstruct K6 from the underlying NHIS items.

Pornography search interest. We use the Google Trends monthly U.S. search-interest index for the literal query “porn,” January 2004 through December 2024 (Google, 2026). The index measures the relative popularity of the query against total U.S. search volume, normalized so the peak month equals 100.

X-rated movie viewing. We use the General Social Survey 1972–2024 Cumulative Data File, Release 3 (NORC at the University of Chicago, 2026), to construct the share of GSS respondents aged 18–44 who reported viewing an X-rated movie in the past 12 months.

4 Descriptive evidence

Figure 2 shows substantial spatial variation in AT&T mobile broadband coverage as of December 31, 2010, spanning the roughly five-year period from June 2007 through February 2011 during which the iPhone was exclusively available on AT&T cellular plans. Coverage varies widely both across and within states. The accompanying bar charts summarize changes in birth rates by age group across three categories of counties: those with little or no AT&T coverage (less than 10% of the population covered), those with partial coverage (10–90%), and those with near-universal coverage (more than 90%).

Simple comparisons of mean changes in birth rates across these coverage cate-

gories reveal large and systematic differences during the AT&T exclusivity period. Teen births declined by 13.8% in counties without AT&T coverage, compared to declines of 18.9% in counties with partial coverage and 26.0% in counties with near-universal coverage. Similar patterns appear at older ages. Births to women in their twenties fell by 10.0% in counties without coverage but by 14.6% in counties with extensive coverage. Among women in their thirties, births rose by 3.8% in counties without AT&T coverage but fell by 1.2% in counties with extensive coverage.⁹

Taken at face value, these comparisons imply a simple difference-in-differences point estimate of -12.2 percentage points among teens, -4.6 for women in their twenties, and -5.1 for women in their thirties relative to women in counties without AT&T coverage. But a causal interpretation of this comparison would rest on a parallel-trends assumption: that high- and low-coverage counties would have evolved similarly between 2007 and 2011 in the absence of the iPhone. Given that treated counties are systematically more urban than control counties (Table 1), any other forces causing urban fertility to decline relatively more than rural fertility over this period could generate the same pattern. The empirical strategies that follow address this concern.

5 Empirical strategy

We adopt two identification strategies. The first is an entropy-balanced Poisson event study that reweights control counties to match treated counties on observable demographics, then estimates the year-by-year effect of the AT&T treatment on age-specific births using a Poisson specification with two-way fixed effects. The second is a synthetic difference-in-differences (SDID) estimator that reweights both

⁹For each age group, the difference in the change in birth rates between counties without and with extensive AT&T coverage is statistically significant at the 1% level.

control counties and pre-treatment years to align the reweighted-control outcome trajectory with the treated trajectory before the iPhone’s 2007 introduction. The two designs address the urbanicity-driven parallel-trends concern raised in [Section 4](#) through different mechanisms: entropy balancing matches the treated and reweighted-control means on observable covariates—neutralizing differential trends to the extent that level imbalance on these characteristics drives them—while SDID aligns the reweighted-control outcome trajectory with the treated trajectory in the pre-period. Both estimates likely understate the iPhone’s effect, since control-county residents could partially access iPhone-mediated content through Wi-Fi or rival-carrier handsets. We describe each in turn below.

5.1 Entropy-balanced Poisson event study

[Table 1](#) documents that treated counties (those with >90% AT&T 3G coverage) are substantially more urban, White, Republican-leaning, and affluent than control counties. To address this imbalance, we apply the entropy-balancing reweighting of Hainmueller (2012), which solves for the entropy-minimizing set of control-county weights that equalize the treated and reweighted-control means of a specified set of covariates. We prefer entropy balancing to propensity-score matching following King and Nielsen (2019), who show that the latter can increase rather than reduce model dependence.

We first divide counties into six strata defined by the interaction of a binary urban indicator (above versus below 33% urban population) and three Census region groupings (Northeast and Midwest, South, and West). Within each stratum, we then balance on four county-level characteristics—the Black population share, the Hispanic population share, the urban-population share, and the 2008 Republican

presidential vote share as a measure of county political orientation.¹⁰ After balancing within strata, we apply a stratum reweight that equalizes each stratum’s share of the balanced-control sample with its share of the treated sample, producing a balanced control pool whose marginal means of the four covariates match the treated marginal means by construction. This two-step procedure is a provides a post-stratification adjustment in the spirit of Deming and Stephan (1940) designed to isolate the sources of identifying variation within urban and region strata. The right-hand columns of [Table 1](#) reports summary statistics for this entropy-balanced control group: the urban-population share rises from 28.2% in the unbalanced controls to match the 66.4% treated mean, and the imbalances in the Black share, Hispanic share, and Republican vote share are likewise eliminated. Balance comes at a cost: equalizing the marginal means requires putting high weight on a small number of treated-like controls. The Kish (1965) effective sample size of the balanced control pool is 77 out of 1,399 raw controls ([Table 1](#)), a constraint we return to in [Section 6](#).

We balance on compositional features of counties rather than on median household income or housing prices, both of which are themselves time-varying outcomes over the 2007–2011 window—a period that overlaps with the Great Recession and the housing crash. Time-varying economic confounders enter our specifications through the regression control vector X_{ct} described below—annual unemployment, the poverty rate, median household income, and the year-over-year change in the Federal Housing Finance Agency’s county Home Price Index—rather than through the balancing weights, and our final specification additionally absorbs state-level economic shocks through state-by-year fixed effects.

On the balanced sample, we estimate, separately for each five-year age band, the

¹⁰The three demographic shares are measured in 2006. We pool the Northeast and Midwest into a single region grouping to ensure sufficient cell counts, as each individually has relatively few counties at the >90%/<10% AT&T cutoff.

Poisson event study

$$\mathbb{E}[B_{ct} | \cdot] = \exp\left(\alpha_c + \gamma_t + \sum_{\tau \neq 2007} \beta_\tau D_c \cdot \mathbf{1}\{t = \tau\} + X'_{ct} \delta\right) \cdot \text{Pop}_{ct}, \quad (1)$$

where B_{ct} is births to women in the relevant age band in county c and year t ; $D_c \in \{0, 1\}$ is the binary AT&T treatment indicator; Pop_{ct} is the corresponding female population, included as an exposure offset; α_c and γ_t are county and year fixed effects; and X_{ct} is a vector of time-varying controls. The event-time coefficients β_τ are normalized to zero in 2007, the last pre-iPhone birth year given gestation timing. We estimate the post-gestation ATT as the mean of the 2008–2011 event-time coefficients with a delta-method standard error derived from the county-clustered variance–covariance matrix of $\hat{\beta}$.¹¹

We report estimates under four nested specifications, all on the entropy-balanced sample: (1) no time-varying covariates; (2) Black, Hispanic, and other-race shares interacted with year (non-Hispanic White omitted as the reference category); (3) + county-level economic conditions (annual unemployment, the poverty rate, median household income, and the year-over-year change in the FHFA Home Price Index) and state-level reproductive- and welfare-policy exposures (parental-involvement laws, mandatory waiting periods, over-the-counter emergency contraception, the contraceptive insurance mandate, the family-planning Medicaid expansion, an active family-cap indicator, and the real maximum monthly TANF benefit); and (4) + state-by-year fixed effects (which absorb all state-level policy controls due to collinearity).¹² Unbalanced versions of these specifications—dropping the entropy-

¹¹Standard errors treat the entropy-balance weights as fixed in the Poisson estimation, following Hainmueller (2012, §3.4).

¹²Time-varying economic controls in Specifications 3 and 4 (unemployment, poverty, median income, HPI YoY change) could be bad controls if iPhone adoption causally affects local economic conditions over 2007–2011. The stability of estimates across Specifications 1–4 is informal evidence that the economic controls are not absorbing the treatment, though we cannot rule the channel out a priori.

balance weights—and OLS analogues are reported in [Figure A.3](#).

5.2 Synthetic difference-in-differences

The SDID estimator of Arkhangelsky et al. (2021) extends standard two-way fixed-effects difference-in-differences by reweighting both control counties and pre-treatment years so that the reweighted-control pre-treatment trajectory is approximately parallel to the treated trajectory. The estimator chooses non-negative unit weights $\hat{\omega}_i$ on control counties, together with a free intercept $\hat{\omega}_0$ that absorbs level differences, so that the $\hat{\omega}$ -weighted pre-treatment outcome trajectory of the controls is approximately parallel to that of the treated counties—it is the slope of the trajectory, not its level, that the weights match.¹³ It chooses non-negative time weights $\hat{\lambda}_t$ on pre-treatment years so that, for each control county, the $\hat{\lambda}$ -weighted average of pre-treatment outcomes approximates that county’s post-treatment mean up to a constant. The treatment-effect estimate is then obtained by weighted least squares with cell weights $\hat{\omega}_i\hat{\lambda}_t$ and unit and time fixed effects.

We estimate the SDID separately for each five-year age band on the balanced 2003–2011 panel of log birth rates, restricted to the >90%/<10% AT&T county sample (914 treated, 1,399 control). Treatment begins in 2008, the first birth cohort whose conceptions could be exposed to the iPhone. We implement the estimator via the Stata `sdid_event` package (Ciccia, 2024), an event-study extension of the parent `sdid` command (Clarke et al., 2024), which extends AAHIW’s pooled ATT to event-time coefficients and produces confidence intervals via a county-clustered bootstrap with 500 replications. The post-gestation ATT is the mean of the 2008–2011 event-time coefficients, with the standard error of this mean computed from

¹³We use the default ridge regularization on unit weights from Arkhangelsky et al. (2021, eq. (4)–(5)); the time-weight optimization is unregularized in AAHIW’s main specification (eq. 6), with only a numerically trivial uniqueness penalty ($\zeta = 10^{-6}\hat{\sigma}$) added in AAHIW’s footnote 3 and inherited by the Stata implementation (Clarke et al., 2024).

the bootstrap variance–covariance matrix.

Although AAHIW’s SDID can accommodate time-varying covariates via residualization (Clarke et al., 2024), we present the headline SDID without covariates. SDID’s unit and time weights mechanically absorb the cross-unit confounders that the parallel-trends Poisson must control for explicitly; both designs nonetheless produce stable estimates across all four nested specifications. To demonstrate that the result is invariant to richer specifications, [Table 2](#) reports post-gestation ATTs under the same four nested specifications we adopted for the entropy-balanced Poisson model, and [Appendix Figures A.1](#) and [A.2](#) visualize the full event-time results for each of those specifications.

5.3 Alternative specifications

The estimates from the two designs above are supported by a series of robustness checks.

[Appendix A](#) reports the results of alternative specification choices that share the $>90\%/<10\%$ AT&T 3G coverage definition of treatment and control counties. [Figure A.3](#) compares the results to nine alternatives—three unbalanced Poisson models and six OLS event studies, each unbalanced and entropy-balanced—on a single age-by-age plot. [Figure A.4](#) relaxes the binary bucketing entirely and estimates a Poisson event study using continuous AT&T coverage share and the full county universe, including those counties with 10–90% AT&T 3G coverage that are excluded from analyses of binary treatment. [Figure A.5](#) estimates the SDID with the outcome in levels (births per 1,000 women) rather than logs. And [Figure A.6](#) estimates SDID modelst using a looser cutoff ($>80\%/<20\%$ AT&T) and a tighter cutoff ($>95\%/<05\%$ AT&T).

A second set of checks addresses a concern for the late post-period: that control

counties could acquire smartphones through a non-iPhone Android channel, attenuating the treated–control contrast. The relevant launches are the T-Mobile HTC Dream / G1 on October 22, 2008 (the first Android phone, and T-Mobile-exclusive until late 2009), the Sprint HTC Hero on October 11, 2009, and the Verizon Motorola Droid on November 6, 2009; AT&T did not gain an Android handset until March 2010. [Figure A.7](#) re-estimates the SDID after additionally restricting both arms to counties with less than 10% T-Mobile 3G coverage, eliminating the only pre-2009 non-iPhone smartphone option from the sample. [Figure A.8](#) truncates the panel to 2003–2009, so the post-period (2008–2009) excludes the birth years materially exposed to the Sprint Hero and Verizon Droid given the ~ 9 -month gestation lag.

Appendix B reports three placebo tests of the SDID design. [Figure B.1](#) estimates a placebo-in-time of SDID on the pre-iPhone sample with the treatment date reassigned to 2003, 2004, 2005, or 2006; a genuine 2008 effect should leave these falsified dates near zero. [Figure B.4](#) and [Figure B.6](#) apply the SDID design to counties served by Verizon and Sprint, holding fixed the 2008 first-treated-year coding used for AT&T. The mean of the 2008–2009 event-time coefficients is a placebo on a 3G-covered carrier that lacked both the iPhone and Android during that window; a null estimate confirms the design is not capturing a generic 3G-coverage advantage. The 2011 coefficient on each cross-carrier placebo is *not* a placebo but a direct SDID estimate of the rival-carrier Android channel arriving on schedule, with results that corroborate the broader smartphone reading.

6 Results

6.1 Entropy-balanced Poisson

Figure 3 reports the entropy-balanced Poisson event study by five-year age band under Specification 4 (state-by-year fixed effects). The results of all four specifications, ranging from no controls to full controls and state:year fixed effects, are reported in Table 2 and visualized in Appendix Figure A.1. The estimates are similar across specifications; the ATT averaged over 2008–2011 under Specification 4 is reported in the panel subtitle of Figure 3.

The estimates suggest large negative effects on births to young women. At ages 15–19, log birth rates in AT&T-treated counties fall sharply below controls beginning in 2008, with an average ATT of -0.046 , or a 4.5% decline in the age-specific birth rate ($p < 0.01$).¹⁴ At ages 20–24 the corresponding ATT is -3.2% ($p < 0.01$). The 25–29 panel shows a smaller effect of -1.0% ($p = 0.36$). At ages 30–34 the ATT is -2.7% ($p < 0.01$); at 35–39 it is -1.4% ($p = 0.23$); and at 40–44 it is -6.4% ($p = 0.16$). Pre-period coefficients are generally not statistically significant, consistent with the balanced-Poisson identifying assumption that, conditional on the balanced-control sample and the time-varying controls, treated and control counties’ birth rates would have evolved similarly absent the iPhone. However, we note that they are also somewhat imprecise and there is visual evidence that births to the younger age groups begin to decline a year prior to iPhone introduction. This pre-trend is substantially smaller than the one visible in the unbalanced continuous-coverage Poisson reported in Figure A.4, which produces ATTs of comparable magnitude to the balanced sample, but more pronounced pre-period divergence between treated and control trajectories.

¹⁴Throughout the paper, we translate log-rate coefficients into exact percent changes using $100 \cdot (\exp(\hat{\beta}) - 1)$.

This lack of precision of the balanced estimates likely reflects the fact that they rest on a small effective sample. The entropy-balancing weights required to match the treated counties’ demographic profile—most consequentially, the shift from a 28.2% to a 66.4% urban population share—concentrate weight on the handful of control counties with large urban populations—the rare controls whose covariate profile resembles a typical treated county. The Kish (1965) effective sample size for the balanced control pool is 77 counties, against the 1,399 raw controls reported in [Table 1](#). The balanced estimates are therefore identified from a narrow slice of the control universe, and their stability across specifications notwithstanding, the precision of the design is constrained by that effective sample. Beyond precision, the small effective sample also makes the point estimate sensitive to the idiosyncratic pre-trajectories of these particular 77 controls—a fragility the SDID estimator addresses through a different mechanism. We turn next to the synthetic difference-in-differences estimator, which addresses the same urbanicity-driven imbalance through outcome-trajectory matching rather than covariate-level matching.

6.2 Synthetic difference-in-differences

[Figure 4](#) reports the SDID event study by five-year age band. The SDID estimates are generally larger in magnitude than the balanced-Poisson results. At ages 15–19 the average ATT is -8.0% . At ages 20–24 the ATT is -6.6% ($p < 0.01$); at 25–29 it is -1.3% ($p < 0.01$); at 30–34 it is -3.6% ($p < 0.01$); at 35–39 it is -1.3% ($p = 0.04$); and at 40–44 it is -2.5% ($p = 0.03$).

The balanced Poisson and SDID estimates for young women differ by roughly a factor of two. The gap traces back to how each estimator handles the pre-iPhone period: entropy balancing matches the 2006 levels of selected covariates; SDID matches the pre-period outcome trajectory. The entropy-balanced Poisson event study ([Fig-](#)

ure 3) shows treated counties' birth rates sitting modestly above balanced controls' in the early pre-period, with a decline toward the controls' level beginning in 2006—a year before the iPhone. The Poisson takes this pre-period as observed and measures the post-2007 ATT against the 2007 baseline. SDID's unit and time weights instead absorb the pre-2007 decline into the no-iPhone counterfactual, producing a larger ATT in absolute value. In our data this difference in pre-period handling is large enough to roughly double the magnitude of the estimated effect at the young ages.

The two estimates bracket the iPhone effect under different assumptions about the 2003–2006 trajectory. If that trajectory reflects a pre-iPhone process that had run its course by 2007, the balanced Poisson is the cleaner number. If it reflects differential dynamics that the balancing covariates do not absorb, SDID's reweighting better captures the drift. The placebo-in-time tests in [Figure B.1](#) and [Figure B.2](#) are largely null under both designs, with one notable exception: at the falsified 2006 introduction date for ages 20–24, the entropy-balanced Poisson produces a placebo coefficient of -0.040 ($p < 0.01$)—in the same direction as the post-2007 estimate and roughly equal in magnitude to the SDID–Poisson gap at that age. This is consistent with the entropy-balanced Poisson failing to absorb the 2006 pre-iPhone decline noted in [Figure 3](#) in at least that cell. We continue to report both estimates as a range, but treat the SDID as the more credible central estimate.

6.3 Plausibility

How large is the per-iPhone-owner effect implied by these results? One might be tempted to do the obvious thing: scale the county-level ATT up by the iPhone ownership share to recover a per-owner number. But the exercise is fraught. The iPhone is not a treatment that operates at the individual level. Whether one's own phone matters likely depends on whether one's peers have phones; a phone in a friend

group full of non-owners is a different intervention than a phone in a group where everyone has one. Spillovers run between phone-owning peers and their non-owning friends, and operate at the level of the group, not just the match: if smartphones reduce friend-group meetups and parties, then matches that would have formed under no-iPhone simply never do—the unformed match is itself the outcome. We cannot cleanly recover an individual-level effect from county-level data when the treatment operates through the social structure those data aggregate across.

With that caveat—and treating the exercise as a plausibility assessment—a back-of-the-envelope calculation goes as follows. In the May 2011 Pew Mobile Survey (Smith, 2011), 12.0% of women and 13.1% of men aged 20–24 owned an iPhone nationally, which scales to 14.0% and 15.2% within counties with AT&T 3G service.¹⁵ Call a *match* a pairing of two people who could, in principle, reproduce. Assuming random matching and no change in match probabilities due to the iPhone—both strong assumptions—then probability that at least one partner in a random match owns an iPhone is approximately 27%.¹⁶

A model in which unexposed matches produce births at the no-iPhone counterfactual rate and exposed matches produce births at $(1 - \delta)$ times that rate implies $ATT = \log(1 - \pi\delta)$, which inverts to $\delta = (1 - \exp(ATT))/\pi$. The SDID’s 2011 endpoint at 20–24 of -0.107 log points—a 10.1% county-level reduction—yields $\delta \approx 0.37$, or a 37% reduction in birth rates among iPhone-exposed matches. The entropy-balanced Poisson’s 2011 endpoint at the same age of -0.038 log points—a 3.7% county-level reduction—yields $\delta \approx 0.14$. The two estimators thus bracket the implied effect on iPhone-exposed matches at 14–37%—within the range a partnering-

¹⁵The Pew survey is national. We assume essentially all iPhone owners in the May 2011 sample lived in counties with AT&T 3G coverage—a defensible approximation given the iPhone’s AT&T exclusivity from June 2007 through February 2011, ending just three months before the survey. Rescaling by the inverse of the U.S. population share of such counties (approximately 86%) gives $12.0\%/0.857 \approx 14.0\%$ for women and $13.1\%/0.857 \approx 15.2\%$ for men.

¹⁶Calculated as $1 - (1 - 0.140)(1 - 0.152) = 0.27$.

delay or dating-app-mediation channel could plausibly produce over the four-year AT&T exclusivity window. At 15–19 the same calculation yields a larger $\delta \approx 40$ –56% (using Pew’s 18–19 cell as a proxy for the unobserved 15–17 ownership rate), consistent with teens combining lower iPhone exposure with a larger per-match response. At ages 25–39, the implied δ ranges from roughly 5% to 14%—higher iPhone ownership paired with a smaller behavioral response per exposed match.

6.4 Heterogeneous effects

Figure 5 reports the post-gestation SDID ATT (mean of 2008–2011 event-time coefficients on the log birth rate) across population subgroups. The first marker in this figure is the estimated ATT for the pooled sample representing births to all women aged 15–44: a 4.8% ($p < 0.01$) decline. Each subsequent marker is a separate SDID estimate on the corresponding subsample, presenting results by 5-year age groupings (corresponding to the models presented in Figure 4), by race and ethnicity, parity, marital status, and education. As a whole, the results illustrate that the diffusion of the iPhone had the largest estimated effects on births to young women aged 15–24 (declines of 8.0% at 15–19 and 6.6% at 20–24, against 1.3 to 3.6% at the older age groups). The decline is broadly similar across White (–4.2%, $p < 0.01$) and Hispanic (–4.9%, $p < 0.01$) women but absent among Black women, where the point estimate is a small positive and statistically indistinguishable from zero (+1.0%, $p = 0.31$). The null Black-women estimate is harder to reconcile with the broad smartphone-diffusion story. Possible explanations include differential iPhone-vs.-Android adoption within the smartphone-using population, and race-specific pre-2007 birth-rate dynamics that the SDID’s outcome-trajectory matching may not neutralize at the subgroup level. There is no evidence of differential effects on births by parity (first births: –4.3%; subsequent: –4.4%; both $p < 0.01$) or marital status (married:

−3.9%; unmarried: −4.3%; both $p < 0.01$). The estimated effects are larger on women with some college education but not a completed degree (−5.8%, vs. −4.3%, −3.2%, and −2.1% for less than high school, high school, and college-plus respectively), but this may simply reflect that women aged 18–24 who attend college are likely to still be in school.

Figure A.9 reports the analogous decomposition under the entropy-balanced Poisson Specification 4. The estimates are noticeably less precise—consistent with the small effective control sample noted earlier—but show qualitatively similar patterns, with larger effects on young women and on White and Hispanic women.

6.5 Robustness to alternative specifications and placebo tests

The headline results survive an extensive set of robustness checks reported in Appendix A and a battery of placebo tests reported in Appendix B (designs described in Section 5). We summarize the outcomes briefly here.

The post-gestation ATTs at the 15–19 and 20–24 ages—where our central effects are estimated—are stable across the twenty alternative specifications reported in Figure A.3: five estimators (SDID, entropy-balanced Poisson, unweighted Poisson, entropy-balanced OLS, and population-weighted OLS) crossed with four nested control specifications, all on the same binary >90/<10 sample. The estimates from these alternative methods cluster tightly around the SDID point estimate at each age. The continuous-coverage Poisson event study on the full county universe in Figure A.4 likewise produces effects of the same sign, dynamics, and significance, indicating the result is not an artifact of the binary bucketing. The level-outcome SDID in Figure A.5 produces post-gestation ATTs of −2.63 ($p < 0.01$), −5.71 ($p < 0.01$), and −3.31 ($p < 0.01$) births per 1,000 women at ages 15–19, 20–24, and 30–34, paralleling the log-rate estimates in implied proportional magnitude. The

threshold sensitivity in [Figure A.6](#) confirms that loosening the cutoff to $>80/<20$ or tightening it to $>95/<05$ does not materially change the estimated effect at any age. The full event-time results under each of the four nested control specifications are visualized in Appendix Figures [A.1](#) and [A.2](#).

The Android-substitution checks reported in [Figure A.7](#) and [Figure A.8](#) address the concern that control counties could have obtained smartphones through Verizon or Sprint Android handsets in the late post-period. Restricting both arms to counties with $<10\%$ T-Mobile coverage—purging the only pre-2009 non-iPhone Android channel from both groups—leaves the 15–19 and 20–24 ATTs at -0.048 ($p < 0.01$) and -0.037 ($p < 0.01$), statistically and substantively similar to the headline. Truncating the panel to a narrow 2003–2009 window that excludes the post-2009 Sprint and Verizon Android launches yields 15–19 and 20–24 post-gestation ATTs of -0.060 ($p < 0.01$) and -0.043 ($p < 0.01$), again negative and statistically significant. The narrow-window magnitudes are smaller than the full-window headline (-0.081 and -0.066), consistent with iPhone adoption compounding through 2010–2011.

The placebo tests in Appendix B further reinforce the design. The placebo-in-time tests in [Figure B.1](#) (SDID) and [Figure B.2](#) (entropy-balanced Poisson) reassign the treatment date to 2003, 2004, 2005, or 2006 on the pre-iPhone sample. For the SDID, the falsified dates yield post-gestation ATTs near zero, confirming the design does not manufacture an effect from pre-2007 dynamics. The entropy-balanced Poisson placebos are noisier, with several falsified coefficients statistically distinguishable from zero—consistent with the pre-period reweighting issue that motivates our SDID-centered reading. The Verizon and Sprint cross-carrier placebos in [Figure B.4](#) and [Figure B.6](#) apply the SDID design to counties served by Verizon-only and Sprint-only respectively, holding fixed the 2008 first-treated-year coding. At ages 15–19 and 20–24, the 2008–2009 placebo coefficients are precisely null on both

carriers, supporting an iPhone-specific interpretation of the headline effect during the AT&T-exclusivity window. The 2011 coefficient on each cross-carrier placebo is sharply negative at the teen ages— -0.033 at 15–19 on Verizon, -0.080 ($p < 0.01$) at both 15–19 and 20–24 on Sprint—and lines up precisely with the gestation-adjusted arrival of Verizon and Sprint Android, corroborating the broader smartphone reading of the AT&T effect rather than threatening it.

7 How much of the 2007–2011 decline does the iPhone explain?

Figure 6 uses the results of these analyses to produce adjusted predictions of the log fertility rate in an alternate world in which the iPhone was not introduced in 2007. Each panel plots a five-year age band’s actual national birth rate against two no-iPhone counterfactuals: one built from the SDID-implied iPhone log effect (Figure 4), the other from the entropy-balanced Poisson Specification 4 estimate (Figure 3), with a shaded band tracing the range between them.

To compute the national counterfactual fertility rate, we substitute zero AT&T coverage for the actual AT&T coverage in each county, rescale the SDID or balanced-Poisson coefficient τ_t to each county’s own coverage share, apply the resulting log effect to that county’s observed births, sum to a national no-iPhone births total, and divide by the same population denominator used for the observed national rate.¹⁷

We treat 2007—whose births were conceived before the June 2007 launch—as the

¹⁷Let a_c denote county c ’s AT&T coverage share. We assume the iPhone’s log effect scales linearly in coverage, so that county c ’s effect at year t is $\tau_t \cdot a_c / (\bar{a}_T - \bar{a}_C)$, where $(\bar{a}_T - \bar{a}_C) \approx 0.98$ is the treated–control coverage gap (the denominator rescales τ_t , which corresponds to that 0.98 binary contrast, to a per-unit-coverage rate). County-level counterfactual births are then $\text{cfb}_{c,t} = \text{births}_{c,t} \cdot \exp(-\tau_t \cdot a_c / (\bar{a}_T - \bar{a}_C))$, and the national counterfactual rate is $\text{cf rate}_t = 1000 \cdot \sum_c \text{cfb}_{c,t} / \sum_c \text{pop}_{c,t}$. The observed national rate uses the same denominator, $\text{rate}_t = 1000 \cdot \sum_c \text{births}_{c,t} / \sum_c \text{pop}_{c,t}$, so both are population-weighted aggregations of county-level quantities and the cf–actual gap is the iPhone’s effect on the population-weighted national rate.

pre-iPhone baseline.

To make this concrete, consider the calculation at ages 15–19 (top-left panel of [Figure 6](#)). Births to women aged 15–19 fell from 41.5 per 1,000 in 2007 to 31.2 per 1,000 in 2011, an observed decline of 10.3 births per 1,000 women. Under the SDID estimates, the no-iPhone counterfactual birth rate at 2011 would have been 34.4 per 1,000—meaning that 3.2 of the 10.3-birth decline, or 31%, is attributable to the iPhone. Under the balanced-Poisson estimates, the counterfactual rate would have been 33.4 per 1,000, attributing 21% of the decline to the iPhone. Similarly, we attribute 14–40% of the decline in births at ages 20–24 and 21–25% of the decline in births at ages 25–29 to the iPhone.

At ages 30–34 the picture is qualitatively different. Births fell from 100.6 per 1,000 in 2007 to 96.2 per 1,000 in 2011, but the SDID counterfactual at 2011 is 100.9 per 1,000—slightly *above* the 2007 baseline—and the balanced-Poisson counterfactual is 99.1, just below it. The pre-iPhone trajectory had been clearly upward, with the actual rate rising from 95.7 per 1,000 in 2003 to 100.6 per 1,000 in 2007, consistent with the secular trend toward later childbearing. The iPhone is identified here as offsetting that upward trajectory rather than amplifying a negative one.

Taken together, the per-age results imply two aggregate findings. First, the iPhone is estimated to account for 20–35% of the 2007–2011 decline in births to women aged 15–29. Second, at ages 30–44 the picture inverts: the birth rate rose from 50.4 to 50.8 per 1,000 between 2007 and 2011, an increase of 0.7%, against a counterfactual rise of 3.6–4.3% to 52.2 (balanced-Poisson) or 52.6 (SDID) per 1,000, leaving the realized increase 81–84% smaller than its no-iPhone counterpart.

[Figure 7](#) reports the corresponding counterfactual for the aggregate general fertility rate (GFR, births per 1,000 women aged 15–44), with the SDID and balanced-Poisson counterfactuals aggregated from the six age-specific event studies—each

band’s own estimated effect applied to that band’s births and summed to the GFR, so this overall counterfactual is the exact aggregation of the per-age trajectories in [Figure 6](#). Of an observed 6.2-birth decline per 1,000 women, between 2.1 (balanced-Poisson) and 3.2 (SDID) are attributable to the iPhone—33 to 52% of the 2007–2011 GFR decline.

8 Mechanisms

Our research design identifies the iPhone’s effect on age-specific birth rates but is silent on the behavioral channels through which the effect may operate. We propose three candidates, drawing on the literature reviewed in our introduction. First, the iPhone substitutes for in-person interaction, displacing the peer time in which most sexual encounters occur (Twenge, 2017; Twenge and Park, 2019; Twenge et al., 2019; Haidt, 2024); causal evidence from Facebook-access experiments is consistent (Allcott et al., 2020; Braghieri et al., 2022). Second, the iPhone provides more access to information about contraception and abortion, which can raise use of either, conditional on sex. Third, the iPhone provides more access to pornography, which can substitute for partnered sex. The first and third channels reduce sex; the second raises contraception conditional on sex; all three reduce the unprotected sex that produces unintended pregnancies (Buckles et al., 2025).

Below we describe national trends in measures relevant to each channel. [Figure 8](#) illustrates trends in sexual behavior that any of the three channels would affect: sex in the past month, contraception at last sex, and recent sex without contraception. [Figure 9](#) shows in-person social interaction and psychological well-being, for the isolation channel; [Figure 10](#) shows pornography access, for the substitution channel; the information channel has no direct upstream measure and enters only through the contraception panel of [Figure 8](#). None of these series can be linked to AT&T

coverage at the cell-of-residence resolution our design requires; they are descriptive trends in nationally representative survey data, consistent with the channels we propose but not a test of any of them.

Figure 8 shows the proximate determinants of fertility from the NSFG, 2002 to 2017–2019, by five-year age band. Sex in the past month is flat at 15–19 (24% to 25%) but falls at older ages (59% to 54% at 20–24; 75% to 61% at 25–29; 75% to 65% at 30–34).¹⁸ Contraception at last sex rises at every age, by 10 percentage points at 15–19 (84% to 94%) and up to 10 points at older ages. The integrated outcome—recent sex without contraception—falls 65% at 15–19 (4.6% to 1.6%), 27% at 25–29 (19.2% to 14.1%), and 28% at 30–34 (19.2% to 13.8%). These are the channels through which any behavioral shock translates into birth rates; Ueda et al. (2020) confirm the sex-frequency decline at the older young-adult ages using independent GSS data, with the share of 18–24 men reporting no sex in the past year rising from 19% in 2000–2002 to 31% in 2016–2018. The next two figures consider mechanisms through which the iPhone may have produced the shift.

Figure 9 documents the in-person displacement channel, in which the iPhone substitutes for face-to-face interaction. ATUS minutes per day with friends present (excluding work and class) fall from 141 to 43 at ages 15–19 (a 69% decline), 107 to 49 at 20–24 (54%), and 65 to 34 at 25–29 (47%) between 2003 and 2024. Time alone rises by 26% at 15–19, 66% at 25–29, and 70% at 30–34. The NHIS-measured share of adults with a Kessler-6 score in the serious-distress range rises from 2.6% to 7.4% at ages 20–24 and 2.5% to 5.6% at 25–29 between 2007 and 2024. The

¹⁸Our NSFG-based finding of flat sexual activity at 15–19 contrasts with the more frequently cited declining trend in the school-based Youth Risk Behavior Survey (YRBS). Lindberg et al. (2021) compare the two surveys head-to-head and document inconsistent prevalence and trends, attributing the divergence to differences in question wording (the YRBS does not define “sexual intercourse”), differences in contraceptive question formats, and peer-context social-desirability bias in the school-based YRBS. Santelli et al. (2025) document additional concerns about YRBS data quality, including a collapse in the school response rate from 81% (2011) to 40% (2023) and a 4-fold increase in missing data on sexual experience. We rely on the NSFG.

iPhone is the always-available alternative to in-person time; its social-media apps are engineered to sustain attention; both features displace the peer time that produces sexual encounters. The cohort-level decline in in-person interaction documented by Twenge et al. (2019) (high-school seniors in 2016 reporting about an hour less per day of in-person socializing than seniors in the late 1980s) and the causal estimates of Allcott et al. (2020) and Braghieri et al. (2022) together provide the empirical scaffolding for this channel.

Figure 10 documents the substitution channel, in which the iPhone provides private mobile access to digital sexual substitutes for partnered sex. The Google Trends annual index for the literal search query *porn* more than doubles over our study period, rising from 40 in 2007 to 86 in 2011. The share of GSS respondents 18–44 reporting they watched an X-rated movie in the past year rises at every age 25 and above between 2000 and 2018, by 22 percentage points at 30–34 (30% to 52%) and 13 points at 40–44 (21% to 34%); the 18–24 cells are too small to read trends reliably. The iPhone made on-demand pornography always-available and private; if the device displaces partnered sex by providing a substitute, the substitution should show up in series like these.

None of these national trends can be assigned a causal interpretation within our design. Appendix C reports an exploratory attempt to combine the limited county and CBSA identifiers in the ATUS with state identifiers to create the minimum spatial “cell” that is county (if identified), balance-of-CBSA, or balance-of-state. Estimating AT&T coverage at the ATUS cell-of-residence, we implement an individual-level analysis of the effects of iPhone diffusion on time spent with friends. The results are consistent with the theory that the diffusion of the iPhone reduced in-person time with friends, but small samples and likely high measurement error in larger cells render them imprecise. For instance, we estimate that the diffusion of

the iPhone reduced time spent with friends by 20.5% ($p = 0.23$) among 15–19 year-olds. Given the lack of precision, we treat these results as no more than a suggestive complement to the descriptive trends, not as an additional causal estimate.

9 Discussion and conclusion

The post-2007 collapse in U.S. fertility has drawn attention across the political spectrum. In May 2024 the Biden Council of Economic Advisers issued an issue brief acknowledging concerns that historically low U.S. fertility creates “significant headwinds to economic growth” and to “the fiscal sustainability of public benefit programs.”¹⁹ On the 2024 campaign trail, Donald Trump called for “a new baby boom,”²⁰ and his administration has since made fertility a recurring policy and rhetorical theme, with senior officials variously describing Americans as “under-babied,” the fertility decline as one of “the most pressing long-term problems facing the country,” and calling for “more babies in the United States of America.”²¹

The concern extends beyond the United States. The IMF characterizes below-replacement fertility as a first-order policy challenge.²² South Korea has spent roughly \$270 billion on pronatalist programs since 2006; France has launched a national “demographic rearmament” campaign; Japan moved to roughly double child-related spending after Prime Minister Kishida warned in January 2023 that addressing the falling birth rate was “now or never.” Fertility has largely continued

¹⁹Council of Economic Advisers, “A First-Principles Look at Historically Low U.S. Fertility and its Macroeconomic Implications,” Issue Brief, May 2024.

²⁰Donald Trump 2024 campaign, Agenda47: “A New Quantum Leap to Revolutionize the American Standard of Living,” donaldjtrump.com.

²¹Remarks by Centers for Medicare & Medicaid Services Administrator Mehmet Oz at an Oval Office maternal health event, May 11, 2026; Russell Vought, “Introduction to the Center for Renewing America FY2023 Budget,” December 7, 2022; Remarks by Vice President J.D. Vance at the March for Life, Washington, D.C., January 24, 2025.

²²Bloom, Kuhn & Prettnner (2025), “The Debate over Falling Fertility,” *Finance & Development*, June 2025; Gruss & Noureldin (2025), “Sustaining Growth in an Aging World,” *Finance & Development*, June 2025.

its long-run decline in each of these countries.

This continued decline is, in our reading, informative. The standard pronatalist toolkit—cash transfers, tax credits, subsidized childcare, extended parental leave—is designed to lower the marginal cost of a child, conditional on the formation of a couple positioned to consider having one. Existing reviews of that toolkit reach a consistent verdict: financial incentives do raise fertility, but modestly and at steep cost. Synthesizing roughly two dozen credibly identified studies, Stone (2020) finds that raising the present value of child benefits by 10 percent of household income increases birth rates by only 0.5 to 4.1 percent, with long-run effects smaller still, as part of the short-run response merely advances the timing of births rather than raising completed family size. Within the United States, Kearney et al. (2022) and Kearney and Levine (2025) find that no standard policy, economic, or social variable can account for the post-2007 break, and conclude that the decline reflects a cohort-level shift in priorities whose proximate cause they treat as an open empirical question.

Our results point to one possible answer. The fertility drop is concentrated among young populations and largely operates through declines in unintended births (Buckles et al., 2025), suggesting the operative margin may be less about the cost of raising a child and more about whether the relationships and sexual activity that produce children are forming at all. The observational evidence in Section 8 is consistent with this reading: as modern smartphones diffused, time spent with friends in person and sexual activity fell sharply alongside rising consumption of pornography, a possible substitute for partnered sex. These mechanisms may not be confined to the young: we observe similar trends for older populations and our SDID estimates of the effect of the iPhone on fertility remain negative and statistically significant at every age band through 40–44, implying that fertility at older ages too

would have been higher absent the iPhone.

We do not claim that the iPhone is the sole cause of the post-2007 decline, nor that no policy lever can move the trajectory. But over the 2008–2011 window that our design identifies, our estimates imply that the introduction of the modern smartphone played a sizable role in the decline in U.S. births. The mechanism evidence suggests this operates through the formation of relationships and the time and inclination for partnered intimacy, not the cost of raising children. If so, the policy instruments to which governments have committed the largest sums—cash transfers, tax credits, subsidized childcare, extended parental leave—do not, on their own, address the behavioral shift our estimates suggest is at work.

References

- Adashi, E. Y., D. P. O’Mahony, and I. G. Cohen (2026). Declining US fertility rates: The specter of a demographic cliff. *The American Journal of Medicine*. Commentary, in press; DOI: 10.1016/j.amjmed.2025.10.028.
- Allcott, H., L. Braghieri, S. Eichmeyer, and M. Gentzkow (2020). The welfare effects of social media. *American Economic Review* 110(3), 629–676.
- Apple Inc. (2007a). Apple Chooses Cingular as Exclusive US Carrier for Its Revolutionary iPhone. Press release, January 9, 2007.
- Apple Inc. (2007b). Apple Reinvents the Phone with iPhone. Press release, January 9, 2007.
- Apple Inc. (2007c). AT&T and Apple Announce Simple, Affordable Service Plans for iPhone. Press release, June 26, 2007.
- Apple Inc. (2007d). iPhone Premieres This Friday Night at Apple Retail Stores. Press release, June 28, 2007.
- Apple Inc. (2011). Verizon Wireless & Apple Team Up to Deliver iPhone 4 on Verizon. Press release, January 11, 2011.
- AppleInsider (2010). Apple, AT&T iPhone Exclusivity Lawsuit Granted Class-Action Status. AppleInsider, July 10, 2010.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review* 111(12), 4088–4118.
- AT&T Inc. (2010a). Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2009. Filed with the U.S. Securities and Exchange Commission, February 25, 2010; Accession No. 0000732717-10-000013.
- AT&T Inc. (2010b). AT&T’s First Android-Powered Smartphone Coming March 7. Press release, February 18, 2010.
- AT&T Inc. (2011). Annual Report on Form 10-K for the Fiscal Year Ended December 31, 2010. Filed with the U.S. Securities and Exchange Commission, March 1, 2011; Accession No. 0000732717-11-000014.
- Blewett, L. A., J. A. R. Drew, M. L. King, K. C. W. Williams, D. Backman, A. Chen, and S. Richards (2024). IPUMS Health Surveys: National Health Interview Survey, Version 7.4. [dataset], Minneapolis, MN: IPUMS.
- Braghieri, L., R. Levy, and A. Makarin (2022). Social media and mental health. *American Economic Review* 112(11), 3660–3693.
- Buckles, K., M. Guldi, and L. Schmidt (2025). The great recession’s baby-less recovery: The role of unintended births. *Journal of Human Resources* 60(1), 224–258.
- Bureau of Labor Statistics (BLS) (2025). American Time Use Survey, 2003-2024 multi-year Activity, Who, Respondent, ATUS-CPS, and Summary files. [dataset].

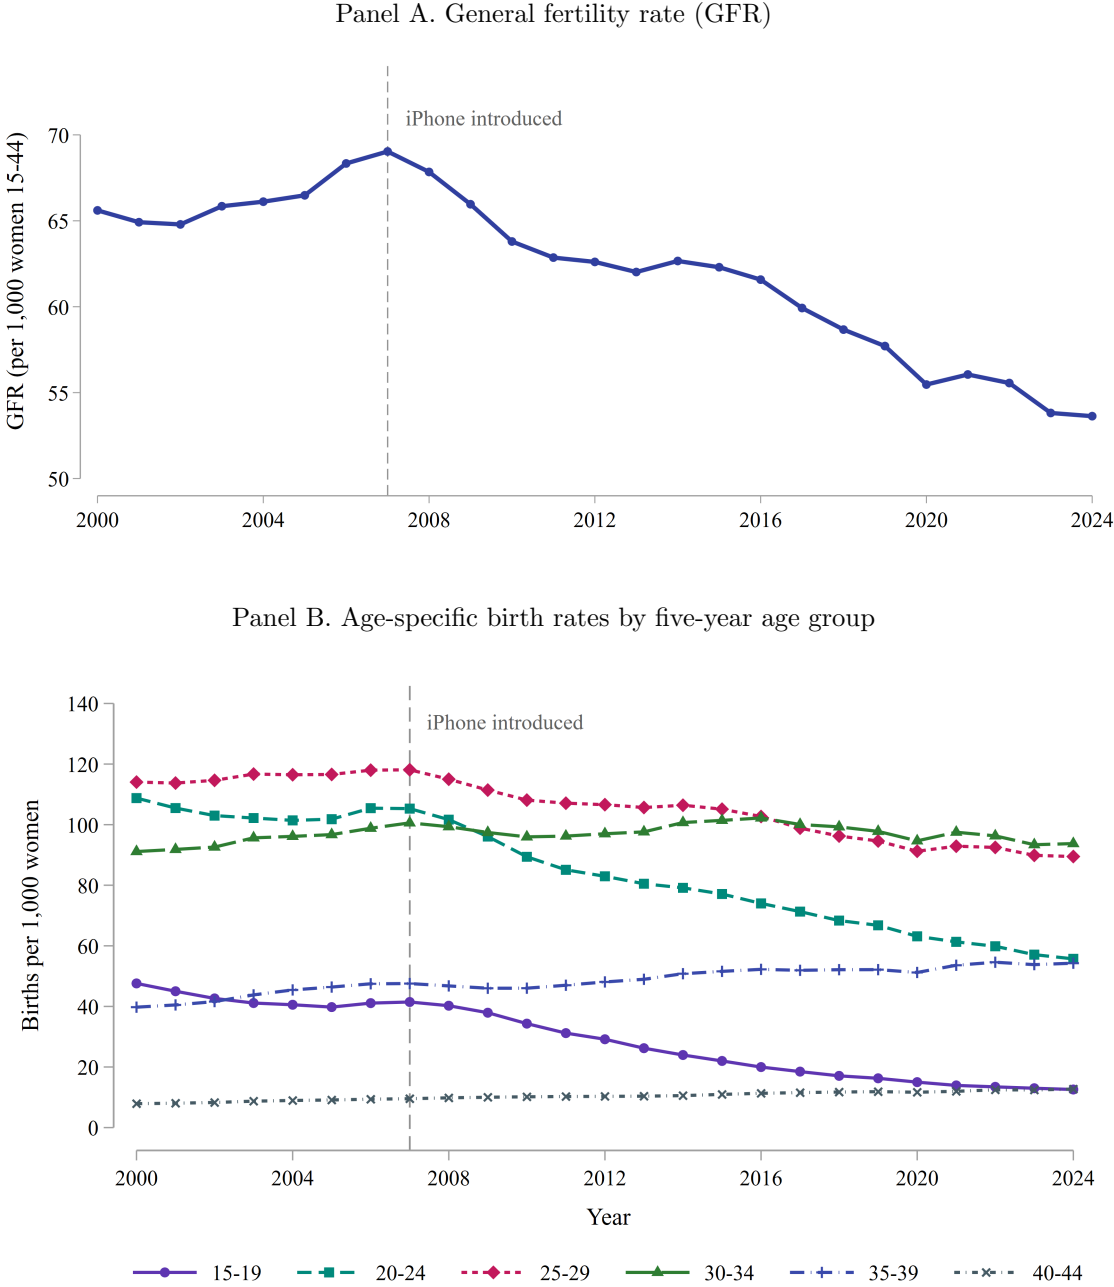
- Bureau of Labor Statistics, Local Area Unemployment Statistics (BLS) (2025). Labor Force data by county and state, 1990-2024. [dataset].
- Chen, B. X. (2009). Rejected by Apple, iPhone Developers Go Underground. *Wired*, August 6, 2009.
- Cherlin, A., E. Cumberworth, S. P. Morgan, and C. Wimer (2013). The effects of the great recession on family structure and fertility. *The Annals of the American Academy of Political and Social Science* 650(1), 214–231.
- Ciccia, D. (2024). A short note on event-study synthetic difference-in-differences estimators. Working paper, arXiv:2407.09565.
- Clarke, D., D. Pailańir, S. Athey, and G. Imbens (2024). On synthetic difference-in-differences and related estimation methods in Stata. *The Stata Journal* 24(4), 557–598.
- CNN Money (2007). Hundreds Camp Out, Crazy for the iPhone. *CNN Money*, June 29, 2007.
- Congressional Budget Office (2014, July). The 2014 long-term budget outlook. Technical Report 45471, Congressional Budget Office.
- Deming, W. E. and F. F. Stephan (1940). On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *Annals of Mathematical Statistics* 11(4), 427–444.
- Detting, L. J. and M. S. Kearney (2023, February). The cyclicity of births and babies’ health, revisited: Evidence from unemployment insurance. Technical Report 30937, National Bureau of Economic Research.
- Federal Communications Commission (2009). Notice of Inquiry on Mobile Wireless Competition. WT Docket No. 09-66, released August 27, 2009.
- Federal Housing Finance Agency (FHFA) (2026). House Price Index Datasets: All-Transactions, county-level, annual, 1975–2025. [dataset]. Vintage: March 31, 2026.
- Ford, G. S. (2021, October). A Quality Check on Form 477 Data: Errors, Subsidies, and Econometrics.
- Google (2026). Google Trends: U.S. monthly search-interest index for the query “porn,” January 2004–December 2024. Google. Accessed 2026.
- Gross, T., J. Lafortune, and C. Low (2014). What happens the morning after? The costs and benefits of expanding access to emergency contraception. *Journal of Policy Analysis and Management* 33(1), 70–93.
- Haidt, J. (2024). *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. New York: Penguin Press.
- Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis* 20(1), 25–46.

- Hudson, N. and H. Moscoso Boedo (2026, April). The collapse of teen fertility in the digital era. Working paper, University of Cincinnati, Carl H. Lindner College of Business. SSRN Working Paper 6676839.
- Kearney, M. S. and P. B. Levine (2009). Subsidized contraception, fertility, and sexual behavior. *Review of Economics and Statistics* 91(1), 137–151.
- Kearney, M. S. and P. B. Levine (2015). Investigating recent trends in the U.S. teen birth rate. *Journal of Health Economics* 41, 15–29.
- Kearney, M. S. and P. B. Levine (2021). Will births in the U.S. rebound? Probably not. Brookings Up Front Brief.
- Kearney, M. S. and P. B. Levine (2022, July). The causes and consequences of declining US fertility. Technical report, Aspen Economic Strategy Group.
- Kearney, M. S. and P. B. Levine (2025, July). Why is fertility so low in high income countries? Technical Report 33989, National Bureau of Economic Research.
- Kearney, M. S., P. B. Levine, and L. Pardue (2022). The puzzle of falling US birth rates since the great recession. *Journal of Economic Perspectives* 36(1), 151–176.
- King, G. and R. Nielsen (2019). Why propensity scores should not be used for matching. *Political Analysis* 27(4), 435–454.
- Kish, L. (1965). *Survey Sampling*. New York: John Wiley & Sons.
- Kruger, L. G. and A. A. Gilroy (2009, February). The Evolving Broadband Infrastructure: Expansion, Applications, and Regulation. Technical Report R40230, Congressional Research Service.
- Lindberg, L. D., R. H. Scott, S. Desai, and Z. H. Pleasure (2021). Comparability of estimates and trends in adolescent sexual and contraceptive behaviors from two national surveys: National survey of family growth and the youth risk behavior survey. *PLOS ONE* 16(7), e0253262.
- MIT Election Data and Science Lab (2018). County Presidential Election Returns 2000–2024. Harvard Dataverse [dataset], V20.
- Myers, C. and D. Ladd (2020). Did parental involvement laws grow teeth? The effects of state restrictions on minors’ access to abortion. *Journal of Health Economics* 71, 102302.
- Myers, C. K. (2021). Cooling off or burdened? The effects of mandatory waiting periods on abortions and births.
- Nasr, A. (2018). The FCC’s Approach to Data Is (Still) a Disaster.
- National Center for Health Statistics (2025). Natality all county files, 2020-2024.
- National Center for Health Statistics (NCHS) (2024). National Survey of Family Growth (NSFG), public-use female respondent files, cycles 2002 through 2017–2019. [dataset].
- National Telecommunications and Information Administration (NTIA) (2010). State Broadband Initiative (SBI) Wireless Data, Current as of June 30, 2010.

- NORC at the University of Chicago (2026). General Social Survey 1972–2024 Cumulative Data File, Release 3 (March 2026). [dataset].
- Patel, N. (2010). Confirmed: Apple and AT&T Signed Five-Year iPhone Exclusivity Deal – But Is It Still Valid? Engadget, May 10, 2010.
- Rachfal, C. L. (2021). Broadband data and mapping: Background and issues for the 117th Congress. Technical Report R45962, Congressional Research Service.
- Santelli, J. S., D. E. Malden, R. A. Moss, M. Finkelstein, and L. D. Lindberg (2025). Age bias, missing data, and declining response rates in the national youth risk behavior survey and their influence on estimates of trends in adolescent sexual experience, 2011–2023. *Sexuality Research and Social Policy*. Online first.
- Schroeder, J., D. Van Riper, S. Manson, K. Knowles, T. Kugler, F. Roberts, and S. Ruggles (2025). IPUMS National Historical Geographic Information System: Version 20.0. [dataset], Minneapolis, MN: IPUMS. Time series table D15 (Persons by Urban/Rural Status), decennial Census 1990–2020; intercensal years interpolated.
- Smith, A. (2011, July). Smartphone adoption and usage. Technical report, Pew Research Center’s Internet & American Life Project.
- Sobotka, T., V. Skirbekk, and D. Philipov (2011). Economic recession and fertility in the developed world. *Population and Development Review* 37(2), 267–306.
- Sprint Nextel Corporation (2009). Sprint Lands HTC Hero: Its First Android Handset, Available October 11. Press release, September 3, 2009.
- Stone, L. (2018, December). Declining fertility in America. Technical report, American Enterprise Institute.
- Stone, L. (2020, March). Pro-natal policies work, but they come with a hefty price tag. Technical report, Institute for Family Studies. .
- Surveillance, Epidemiology, and End Results Program (SEER) (2026). U.S. County Population Data 1990-2024, Vintage 2024: single years of age by race and Hispanic origin. [dataset], National Cancer Institute. Downloaded April 24, 2026.
- T-Mobile USA (2008). T-Mobile Unveils the T-Mobile G1—The First Phone Powered by Android. Press release, September 23, 2008.
- The CPU Shack Museum (2011). Qualcomm for Apple: The iPhone 4 CDMA. February 8, 2011.
- Twenge, J. M. (2017). *iGen: Why Today’s Super-Connected Kids Are Growing Up Less Rebellious, More Tolerant, Less Happy—and Completely Unprepared for Adulthood—and What That Means for the Rest of Us*. New York: Atria Books.
- Twenge, J. M. (2023). *Generations: The Real Differences between Gen Z, Millennials, Gen X, Boomers, and Silents—and What They Mean for America’s Future*. New York: Atria Books.
- Twenge, J. M., G. N. Martin, and B. H. Spitzberg (2019). Trends in U.S. adolescents’ media use, 1976–2016: The rise of digital media, the decline of TV, and the (near) demise of print. *Psychology of Popular Media Culture* 8, 329–345.

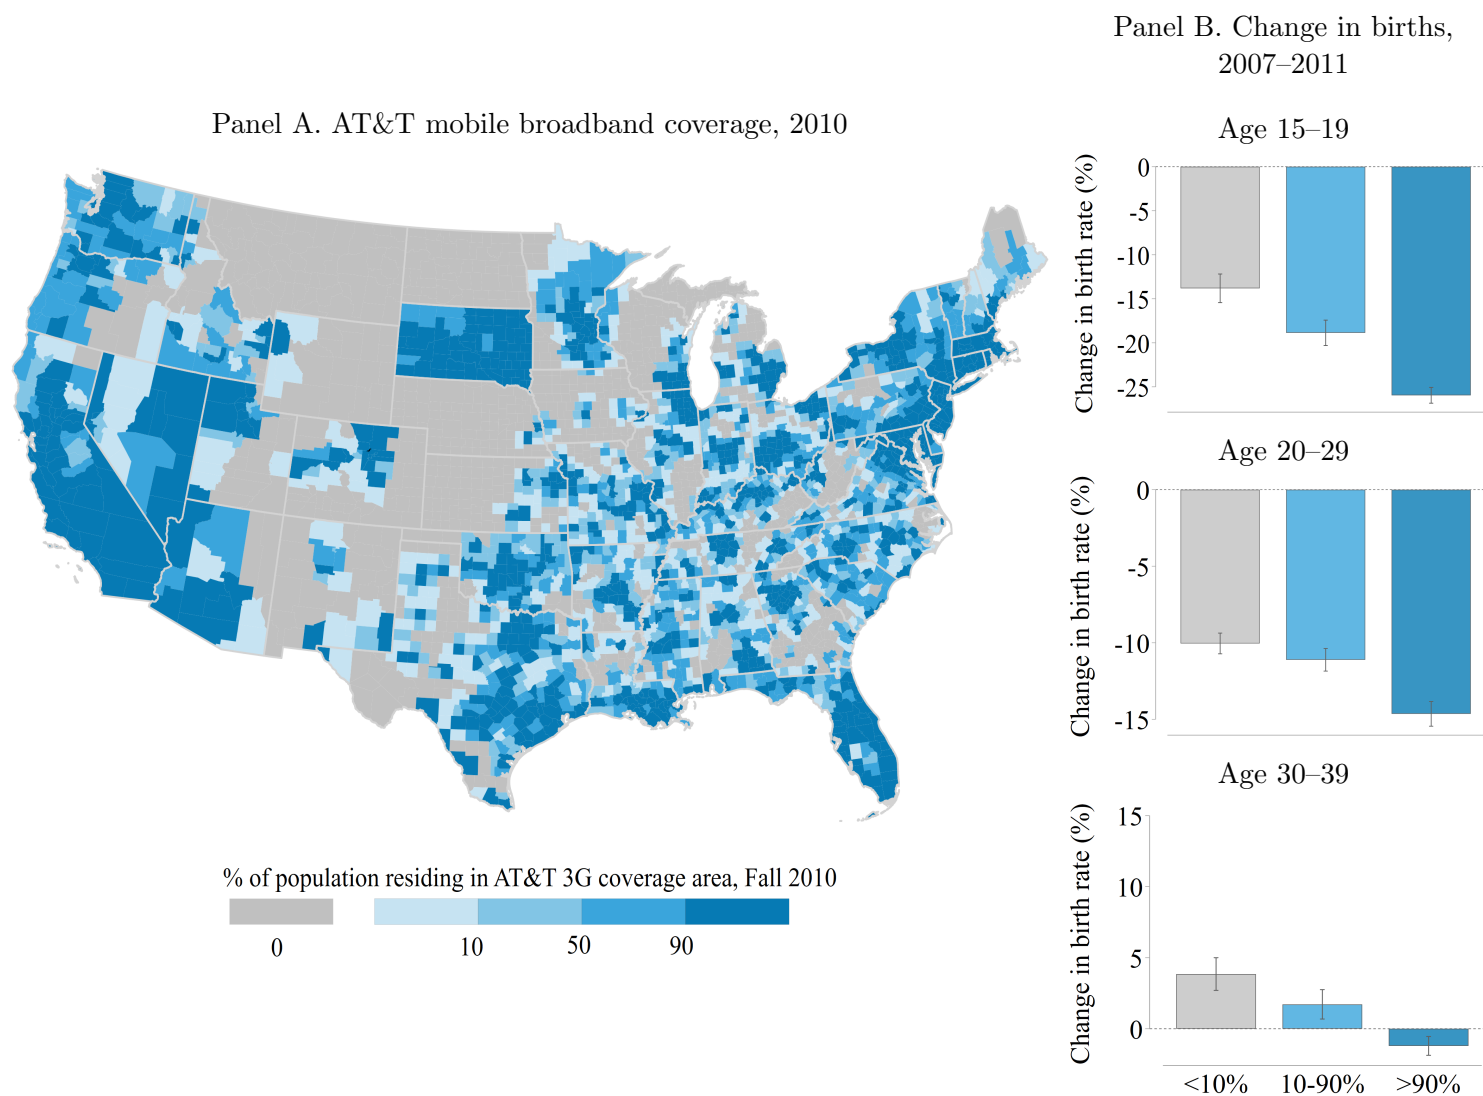
- Twenge, J. M. and H. Park (2019). The decline in adult activities among U.S. adolescents, 1976–2016. *Child Development* 90(2), 638–654.
- Twenge, J. M., R. A. Sherman, and B. E. Wells (2017). Sexual inactivity during young adulthood is more common among U.S. millennials and iGen: Age, period, and cohort effects on having no sexual partners after age 18. *Archives of Sexual Behavior* 46(2), 433–440.
- Twenge, J. M., B. H. Spitzberg, and W. K. Campbell (2019). Less in-person social interaction with peers among U.S. adolescents in the 21st century and links to loneliness. *Journal of Social and Personal Relationships* 36(6), 1892–1913.
- Ueda, P., C. H. Mercer, C. Ghaznavi, and D. Herbenick (2020). Trends in frequency of sexual activity and number of sexual partners among adults aged 18 to 44 years in the US, 2000–2018. *JAMA Network Open* 3(6), e203833.
- Urban Institute (2023). Welfare Rules Database. [dataset]. Maximum monthly TANF benefit for a family of three (Table II.A.4); welfare reform and family cap implementation dates from longitudinal policy tables. Vintage: policies as of July 1, 2023. Accessed 2026.
- U.S. Census Bureau (2024). Small Area Income and Poverty Estimates. [dataset].
- U.S. Senate Committee on Commerce, Science, and Transportation (2009). The Consumer Wireless Experience. Hearing, 111th Cong., 1st Sess., June 17, 2009 (S. Hrg. 111-449).
- U.S. Senate Subcommittee on Antitrust, Competition Policy and Consumer Rights (2009). Cell Phone Text Messaging Rate Increases and the State of Competition in the Wireless Market. Hearing, 111th Cong., 1st Sess., June 16, 2009.
- Verizon Wireless (2009). DROID by Motorola Available Nov. 6 Exclusively at Verizon Wireless. Press release, October 28, 2009.

Figure 1: The post-2007 fertility decline, in aggregate and by age



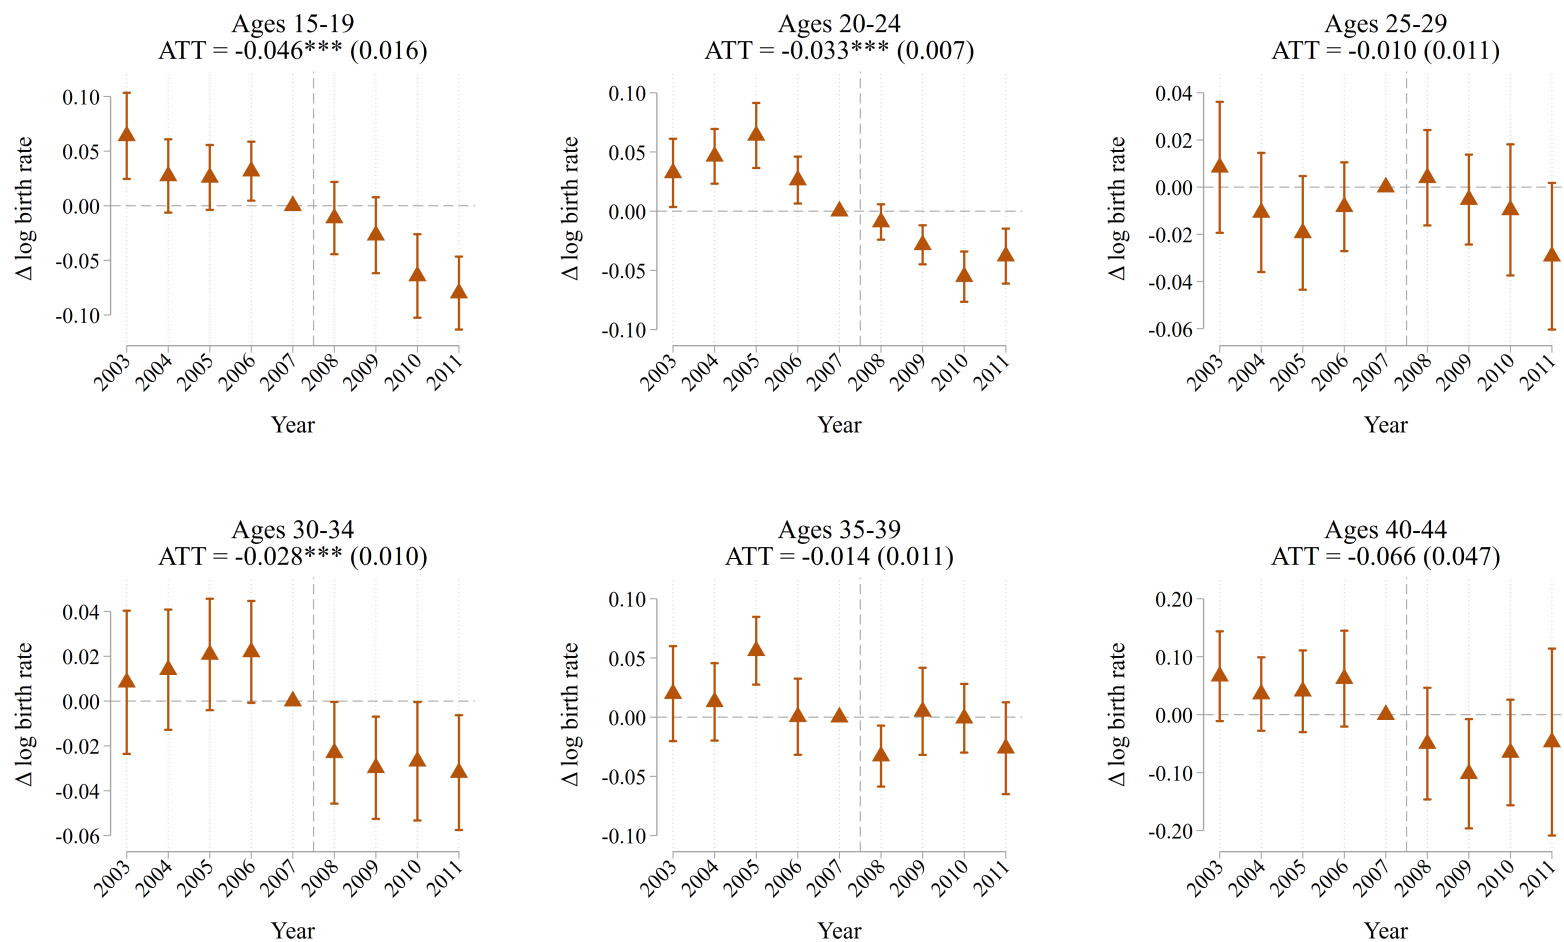
Notes: Panel A plots the U.S. general fertility rate (GFR), the annual number of live births per 1,000 women aged 15–44. Panel B plots national age-specific birth rates (births per 1,000 women in each age group) by five-year age group. Both panels span 2000–2024 and are population-weighted across all counties in the analysis sample. The vertical dashed line in each panel marks 2007, the year the iPhone was introduced. Sources: NCHS (2025), SEER (2026).

Figure 2: AT&T coverage is correlated with relative declines in birth rates



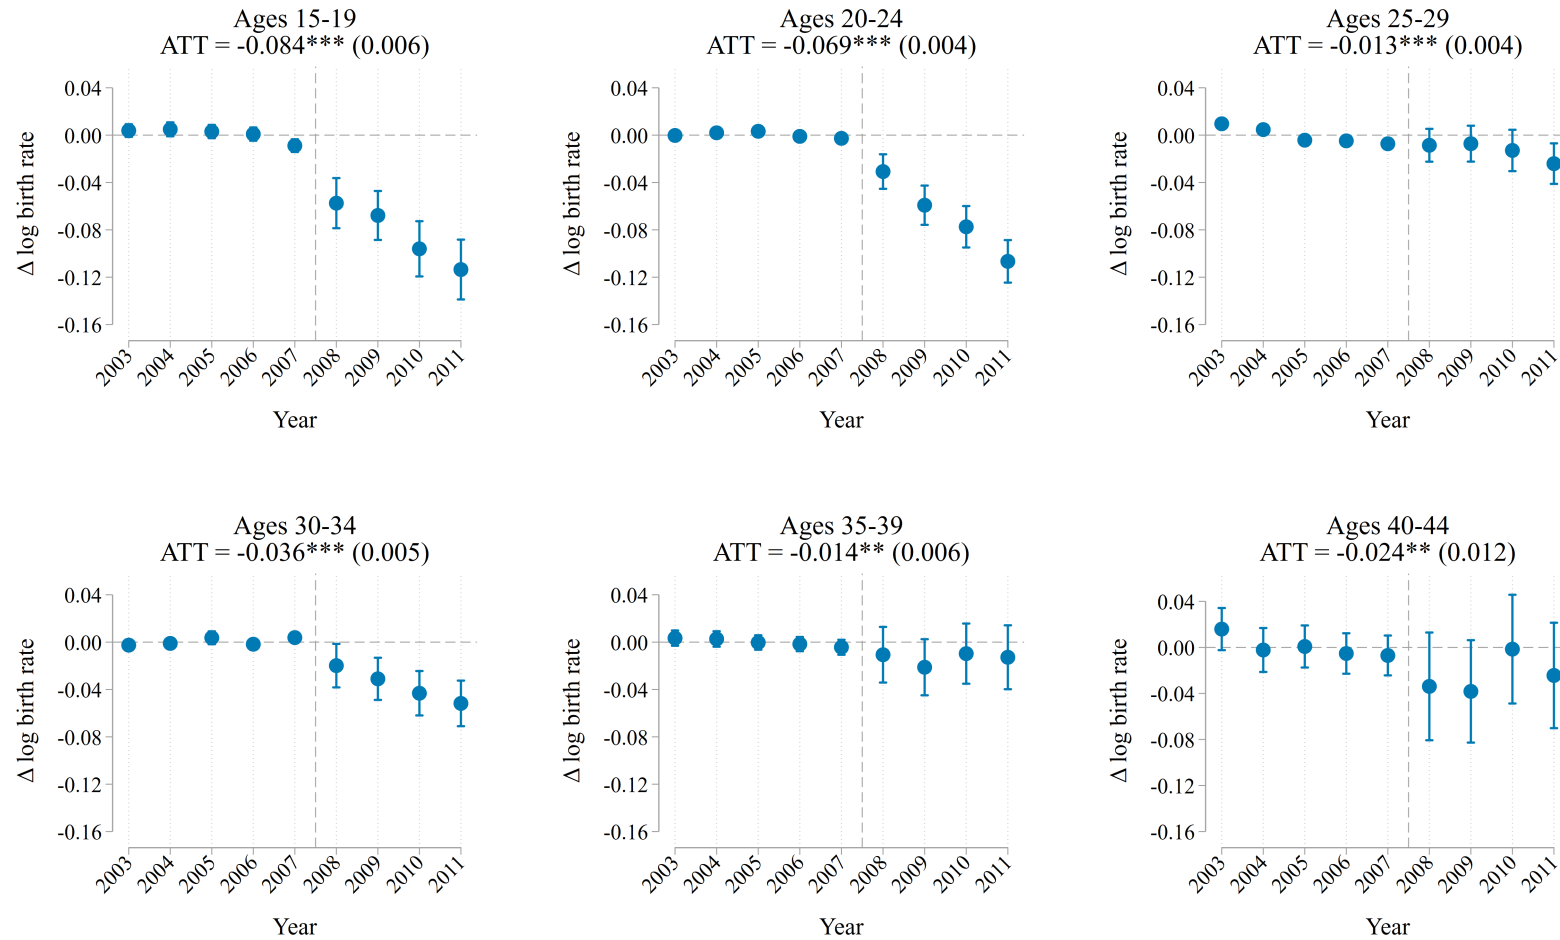
Notes: Birth rates are calculated by the authors using all-county natality files (NCHS, 2025) and aggregations are weighted by the corresponding population. County-level AT&T mobile broadband coverage is constructed from AT&T wireless broadband submissions to the State Broadband Initiative (NTIA, 2010). A census block is coded as covered if SBI reports that AT&T provides wireless broadband service in the block at speeds of at least 200 Kbps, the FCC’s standard definition of broadband at the time. Block-level indicators are averaged unweighted to the block group, and block-group shares are averaged to the county weighted by 2000 Census block-group population, as of December 31, 2010. Bars represent 95% confidence intervals based on t-tests.

Figure 3: Estimated effect of AT&T 3G coverage on birth rates, by age, from an entropy-balanced Poisson event study



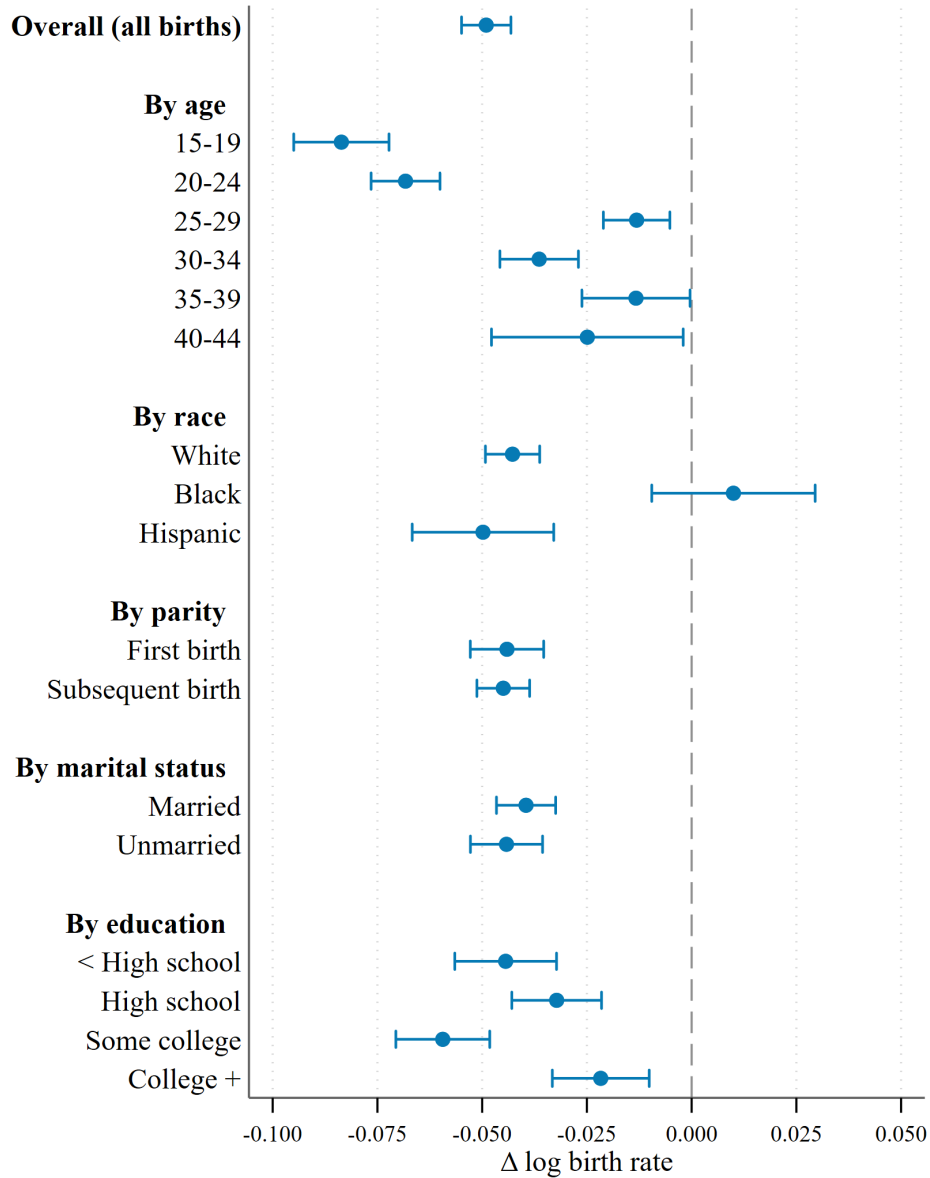
Notes: Each panel reports a Poisson (PPML) event study of age-specific births on the binary AT&T treatment (>90% vs. <10% coverage), with population as the offset. Within six rural/urban \times region (NE+MW, South, West) strata, control counties are reweighted via per-age entropy balancing (Hainmueller, 2012) to match treated counties on Black, Hispanic, Republican-vote, and urban shares; stratum weights are then aligned to the treated distribution. The sample is a balanced 2003–2011 county panel and 2008 is the first treated year (the iPhone launched in June 2007, so 2007 births were conceived pre-iPhone). Specification 4 is shown: county and state \times year fixed effects, demographic shares interacted with year, and time-varying economic controls; state-level policy controls are absorbed by the state \times year fixed effects. See Table 1 for the variables in each group. Four specifications with nested controls—a no-covariate baseline, then adding demographic shares \times year, economic and policy controls, and state \times year fixed effects—are reported in Table 2 and visualized in Appendix Figure A.1. Each panel subtitle reports the post-gestation ATT (mean of the 2008–2011 event-time coefficients). Standard errors are clustered at the county level; bars are 95% confidence intervals.

Figure 4: Estimated effect of AT&T 3G coverage on birth rates, by age, from a synthetic difference-in-differences event study



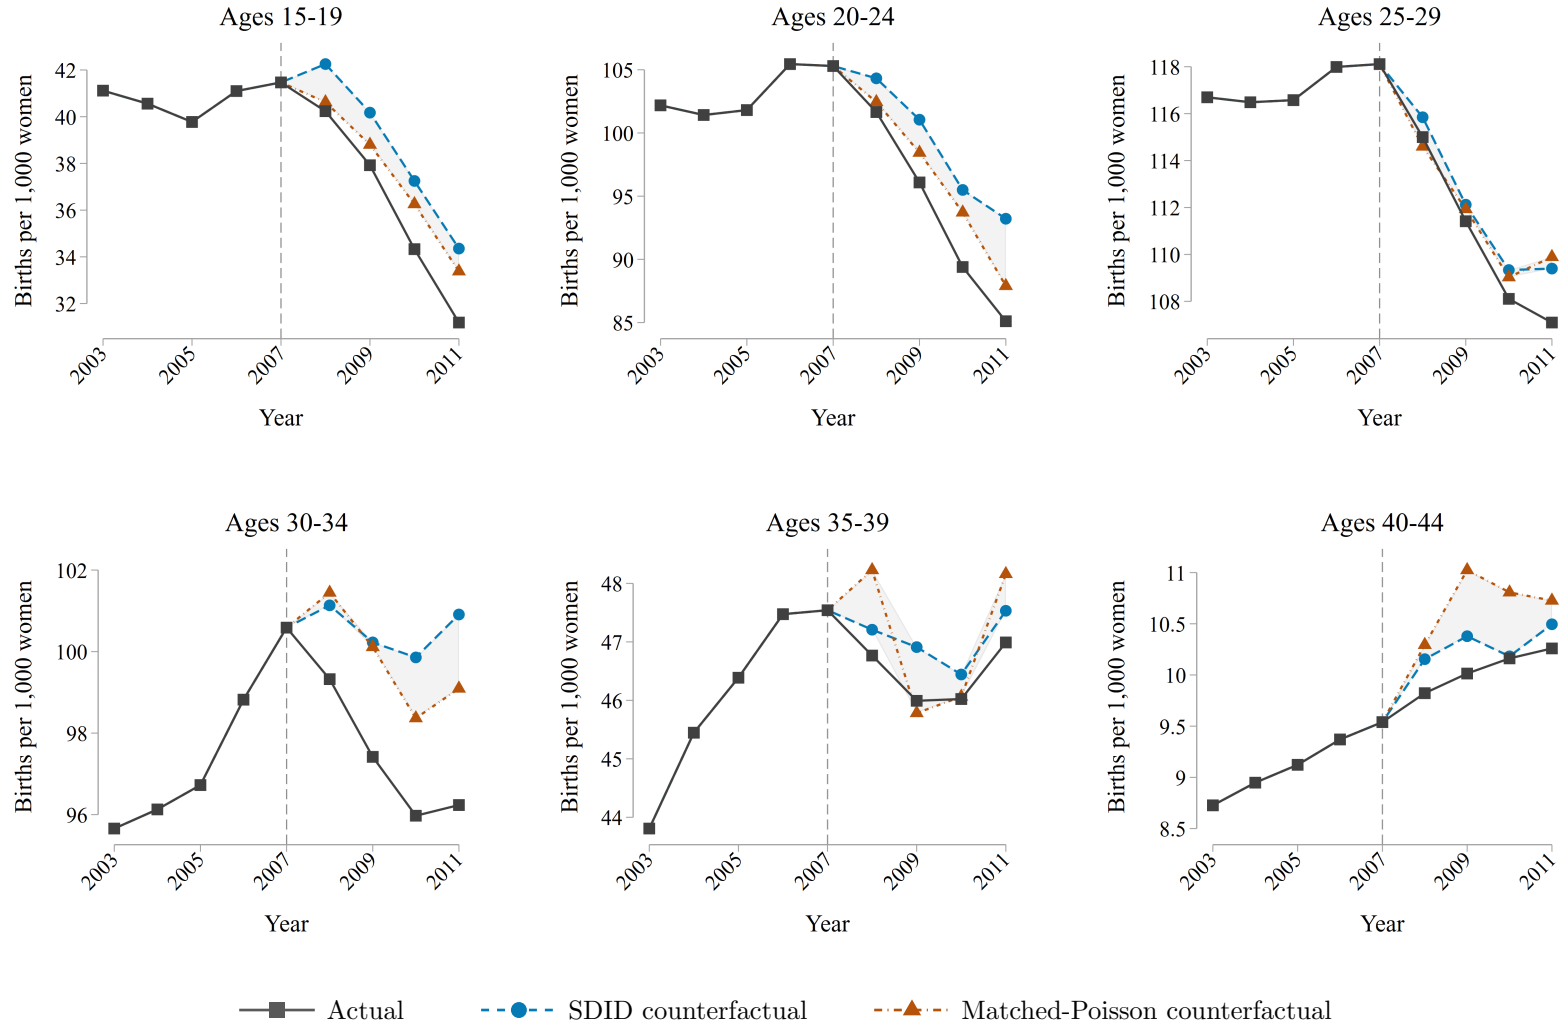
Notes: Each panel reports a synthetic difference-in-differences (SDID) event study of the log age-specific birth rate (Arkhangelsky et al., 2021), with no time-varying covariates. Treated counties have AT&T 3G coverage >90% of the population; control counties <10%; the partially covered middle is excluded. The sample is a balanced 2003–2011 county panel and 2008 is the first treated year (the iPhone launched in June 2007, so 2007 births were conceived pre-iPhone). Four specifications with nested controls—a no-covariate baseline, then adding demographic shares \times year, economic and policy controls, and state \times year fixed effects—are reported in Table 2 and visualized in Appendix Figure A.2. Each panel subtitle reports the post-gestation ATT (the mean of the 2008–2011 event-time coefficients). Inference is by county-clustered bootstrap with 500 replications; bars are 95% confidence intervals.

Figure 5: Estimated effect of AT&T 3G coverage on birth rates, by subgroup, from a synthetic difference-in-differences event study



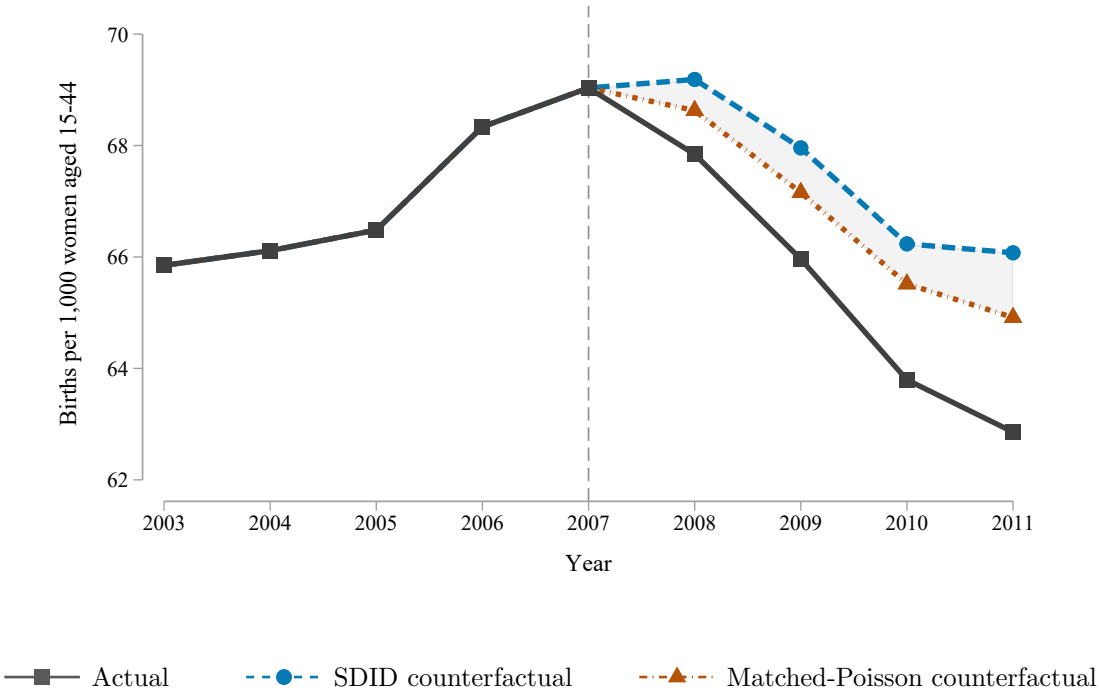
Notes: Each marker is the post-gestation mean of the 2008–2011 event-time coefficients from a synthetic difference-in-differences estimate of the log birth rate, with 95% confidence intervals from a county-clustered bootstrap (500 replications). The six age rows reproduce the by-age estimates in Figure 4. Age- and race-specific rows use group-specific population denominators; the overall, parity, marital-status, and education rows use women aged 15–44 as the denominator. Because NCHS removed the unrevised-certificate maternal-education item beginning in 2011—leaving education unreported for births in states still using the 1989 certificate—the education rows are estimated on the 2003–2010 window; all other rows use the full 2003–2011 window. A parallel entropy-balanced Poisson version of this figure is reported in Figure A.9. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure 6: Actual vs. no-iPhone counterfactual birth-rate trend, by age




Notes: In each panel the solid dark-gray line with square markers is the actual national age-specific birth rate per 1,000 women in the indicated five-year age group. The dashed blue line with circle markers uses the SDID results to estimate a no-iPhone counterfactual and the dash-dot amber line with triangle markers uses the entropy-balanced Poisson to do the same. The shaded gray band traces the range between the two counterfactuals. See text for more detail. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure 7: General fertility rate: actual vs. no-iPhone counterfactual




Notes: The solid dark-gray line with square markers is the actual national general fertility rate (births per 1,000 women aged 15–44), 2003–2011. The dashed blue line with circle markers is the SDID no-iPhone counterfactual (Figure 4) and the dash-dot amber line with triangle markers is the entropy-balanced Poisson Specification 4 counterfactual (Figure 3). The shaded gray band traces the range between the two counterfactuals. The vertical dashed reference marks the June 2007 iPhone launch. See text for more detail. The age-disaggregated trajectory is in Figure 6. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure 8: Trends in sexual behavior



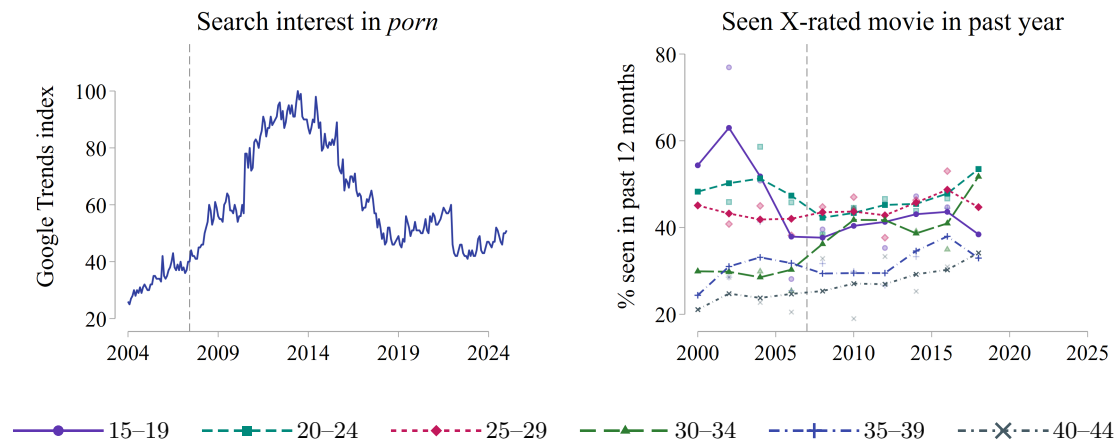
Notes: Each line is a connected by-age series across NSFG cycles 2002 through 2017–2019, plotted at cycle midpoints; women aged 15–44 (pooled $N = 41,500$ across cycles). Left panel is the share who had sex in the past month (interview minus last-sex century-month $\in [0, 1]$; 0 if no recent sex). Middle panel is the share who used any contraception at last sex ($USELSTP = \text{Yes}$), among women who answered the question. Right panel is the integrated proximate outcome: had sex in the past month AND used no contraception at that last encounter; 0 if no recent sex; missing if recent sex but ambiguous $USELSTP$. Dashed line marks the iPhone introduction (June 2007). Survey weights are applied. Source: NSFG (NCHS, 2024).

Figure 9: Trends in time with other people and psychological distress



Notes: Left and middle panels are ATUS minutes per day, lowess-smoothed with faint raw markers behind; sample 15-44, both sexes pooled, 2003-2024 excluding 2020 ($N = 110,323$). Left panel restricts to episodes with a friend (ATUS who-code 54) present, excluding work and class. Middle panel is time alone (who-codes 18, 19). Right panel is the share of NHIS adults 18-44 with $K6 \geq 13$ (serious psychological distress), with K6 reconstructed from the underlying items; NHIS interviews only adults so the youngest bin is 18-19 in the right panel ($N = 22,668$). Dashed line marks the iPhone introduction (June 2007). Survey weights are applied. Sources: ATUS (BLS, 2025); IPUMS NHIS (Blewett et al., 2024).

Figure 10: Trends in Google searches for “porn” and X-rated movie viewing



Notes: Left panel is the Google Trends monthly U.S. search-interest index for the search query “porn,” January 2004 through 2024, normalized so the peak month equals 100; the index measures relative search popularity, not absolute volume. Right panel is the share of GSS respondents aged 18–44 who reported viewing an X-rated movie in the past 12 months, lowess-smoothed with faint raw markers behind, 2000–2024 ($N = 6,945$). GSS minimum interview age is 18, so the youngest bin is labeled 18–19. Dashed line marks the iPhone introduction (June 2007). Survey weights are applied to the GSS series. Sources: Google Trends (Google, 2026); GSS (NORC at the University of Chicago, 2026).

Table 1: Summary statistics by AT&T mobile broadband coverage, 2003–2006

| | Control (<10% AT&T) | Balanced controls | Treated (>90% AT&T) | Excluded (10–90% AT&T) | All counties |
|---|------------------------|----------------------|------------------------|---------------------------|-----------------|
| <i>Mobile broadband coverage</i> (% of women 15–44 residing in covered area, pop-weighted) | | | | | |
| Any national carrier | 87.0 (27.5) | 96.6 (14.4) | 99.8 (0.7) | 95.0 (11.8) | 98.1 (9.8) |
| AT&T | 1.2 (2.4) | 0.4 (1.4) | 99.0 (2.0) | 62.3 (25.1) | 86.5 (30.0) |
| Sprint | 47.2 (39.9) | 82.5 (31.3) | 96.0 (10.8) | 67.9 (31.3) | 88.7 (24.1) |
| T-Mobile | 3.7 (15.5) | 26.6 (36.2) | 83.2 (28.8) | 15.6 (26.5) | 68.9 (40.1) |
| Verizon | 77.2 (39.7) | 88.9 (29.9) | 96.3 (17.9) | 86.7 (30.7) | 93.6 (23.0) |
| <i>Outcome: birth rate</i> (per 1,000 women in age group) | | | | | |
| Total 15–44 | 67.2 (14.4) | 64.8 (17.1) | 66.8 (12.7) | 65.0 (12.2) | 66.5 (13.4) |
| Age 15–19 | 45.7 (25.5) | 48.3 (25.6) | 40.9 (19.9) | 46.3 (21.0) | 44.4 (23.0) |
| Age 20–24 | 149.1 (49.9) | 125.1 (57.5) | 116.7 (44.3) | 138.7 (43.0) | 136.9 (48.6) |
| Age 25–29 | 145.1 (46.9) | 130.9 (34.5) | 130.4 (33.9) | 132.3 (29.1) | 137.5 (39.9) |
| Age 30–34 | 78.2 (31.2) | 76.1 (23.3) | 93.1 (24.3) | 75.4 (20.8) | 81.9 (27.9) |
| Age 35–39 | 29.1 (15.1) | 30.3 (11.3) | 40.2 (15.0) | 29.3 (11.6) | 32.4 (15.1) |
| Age 40–44 | 5.3 (5.2) | 5.2 (3.6) | 7.5 (4.0) | 5.5 (3.7) | 6.0 (4.6) |
| <i>Demographic composition</i> (county population shares, %) | | | | | |
| Non-Hispanic White | 81.2 (21.7) | 75.9 (19.9) | 74.1 (20.0) | 80.7 (19.1) | 79.0 (20.8) |
| Non-Hispanic Black | 8.0 (16.7) | 11.4 (16.0) | 11.3 (14.5) | 10.5 (16.6) | 9.6 (16.1) |
| Hispanic | 8.1 (15.2) | 10.2 (16.1) | 10.0 (13.8) | 6.1 (11.0) | 8.2 (13.9) |
| Other race | 2.7 (7.8) | 2.5 (4.7) | 4.6 (8.1) | 2.7 (5.9) | 3.3 (7.5) |
| <i>Urbanization and economic conditions</i> | | | | | |
| Urban (%) | 28.2 (26.1) | 66.7 (27.7) | 66.4 (28.2) | 33.5 (23.7) | 40.8 (31.0) |
| Unemployment rate (%) | 5.6 (2.0) | 5.3 (1.7) | 5.1 (1.6) | 5.7 (1.7) | 5.5 (1.8) |
| Poverty rate (%) | 15.6 (6.0) | 16.0 (4.9) | 12.6 (5.4) | 14.7 (5.4) | 14.5 (5.8) |
| Med household inc (\$1000) | 56.3 (10.0) | 57.9 (8.9) | 74.0 (19.5) | 61.9 (13.0) | 62.9 (16.0) |
| HPI (% change) | 5.2 (5.4) | 5.5 (5.4) | 7.6 (6.5) | 6.2 (5.5) | 6.3 (5.9) |
| <i>Policy exposure and political environment</i> (state-level policy shares except where noted) | | | | | |
| Parental involvement law (%) | 85.0 (35.3) | 83.5 (36.8) | 72.2 (44.3) | 79.2 (40.1) | 79.8 (39.7) |
| Mandatory waiting period (%) | 11.2 (31.1) | 6.5 (24.3) | 10.4 (29.9) | 13.6 (33.8) | 11.6 (31.5) |
| OTC emergency contraception (%) | 23.8 (39.7) | 24.6 (40.1) | 26.5 (41.5) | 24.2 (40.0) | 24.7 (40.3) |
| Contraceptive insurance mandate (%) | 29.4 (44.8) | 46.8 (49.1) | 37.6 (47.9) | 33.6 (46.8) | 32.9 (46.4) |
| FP Medicaid expansion (%) | 20.9 (39.3) | 20.6 (38.8) | 23.8 (41.5) | 29.2 (44.1) | 23.9 (41.4) |
| Active family-cap policy (%) | 45.1 (49.8) | 45.1 (49.8) | 46.1 (49.9) | 49.2 (50.0) | 46.4 (49.9) |
| Max monthly TANF benefit (\$2024) | 574 (184) | 573 (159) | 627 (245) | 550 (219) | 584 (215) |
| Republican vote share (% county) | 61.5 (12.6) | 55.9 (12.4) | 56.1 (12.7) | 59.7 (11.1) | 59.4 (12.5) |
| Number of counties | 1,399 | 1,399 | 914 | 794 | 3,107 |
| Effective N (Kish) | — | 77 | — | — | — |
| Total female pop 15–44 (millions) | 5.3 | 5.3 | 49.8 | 6.6 | 61.8 |

Notes: County-level means, standard deviations in parentheses, 2003–2006 pre-treatment period. Columns partition counties by AT&T coverage as of December 31, 2010: Control (<10% of population), Treated (>90%), Excluded (10–90%). “Balanced controls” applies entropy-balancing weights (Hainmueller, 2012) within six rural/urban \times region (NE+MW, South, West) strata to match the treated means of the Non-Hispanic Black, Hispanic, Republican-vote, and urban shares, with stratum weights then aligned to the treated distribution; the row reports the Kish (1965) effective N . The mobile broadband coverage panel is population-weighted by women 15–44 (“Any national carrier” approximated by the maximum single-carrier share); all other rows are unweighted county means. Reproductive- and welfare-policy rows are state-by-year exposure measures rescaled to 0–100, except the TANF benefit (real 2024 dollars). Sources: NCHS (2025), SEER (2026), NHGIS (Schroeder et al., 2025), MEDSL (MIT Election Data and Science Lab, 2018), BLS (2025), SAIPE (2024), FHFA (2026), NTIA (2010), Myers and Ladd (2020); Myers (2021); Gross et al. (2014); Kearney and Levine (2009), and the Urban Institute Welfare Rules Database (Urban Institute, 2023).

Table 2: Post-gestation ATT comparison: SDID and entropy-balanced Poisson under nested control specifications

| Age | Specification | SDID | | Matched Poisson | |
|-------|------------------------------|-----------|---------|-----------------|---------|
| | | ATT | SE | ATT | SE |
| 15–19 | No covariates | -0.084*** | (0.006) | -0.071*** | (0.014) |
| | + demographics \times year | -0.082*** | (0.006) | -0.058*** | (0.015) |
| | + econ + policy | -0.070*** | (0.006) | -0.045*** | (0.014) |
| | + state \times year FE | -0.091*** | (0.006) | -0.046*** | (0.016) |
| 20–24 | No covariates | -0.069*** | (0.004) | -0.060*** | (0.012) |
| | + demographics \times year | -0.066*** | (0.004) | -0.044*** | (0.012) |
| | + econ + policy | -0.057*** | (0.004) | -0.037*** | (0.011) |
| | + state \times year FE | -0.048*** | (0.004) | -0.033*** | (0.007) |
| 25–29 | No covariates | -0.013*** | (0.004) | -0.003 | (0.013) |
| | + demographics \times year | -0.012*** | (0.004) | 0.004 | (0.014) |
| | + econ + policy | -0.008** | (0.004) | 0.009 | (0.014) |
| | + state \times year FE | -0.020*** | (0.004) | -0.010 | (0.011) |
| 30–34 | No covariates | -0.036*** | (0.005) | -0.034*** | (0.011) |
| | + demographics \times year | -0.031*** | (0.005) | -0.025** | (0.012) |
| | + econ + policy | -0.026*** | (0.005) | -0.022* | (0.012) |
| | + state \times year FE | -0.040*** | (0.005) | -0.028*** | (0.010) |
| 35–39 | No covariates | -0.014** | (0.006) | 0.013 | (0.016) |
| | + demographics \times year | -0.014** | (0.007) | 0.006 | (0.015) |
| | + econ + policy | -0.008 | (0.007) | 0.012 | (0.015) |
| | + state \times year FE | -0.003 | (0.007) | -0.014 | (0.011) |
| 40–44 | No covariates | -0.024** | (0.012) | -0.039 | (0.061) |
| | + demographics \times year | -0.025** | (0.012) | -0.045 | (0.061) |
| | + econ + policy | -0.019 | (0.012) | -0.041 | (0.060) |
| | + state \times year FE | -0.052*** | (0.012) | -0.066 | (0.047) |

Notes: Post-gestation ATT estimates (mean of 2008–2011 event-time coefficients) on the log age-specific birth rate, by five-year age band, for both estimators under four nested control specifications: (1) no covariates; (2) + demographic shares interacted with year; (3) + economic and policy controls; (4) + state \times year fixed effects, which absorb the state-level policy controls. Table 1 lists the variables in each group. SDID inference is by county-clustered bootstrap with 500 replications. Entropy-balanced Poisson standard errors are clustered at the county level. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The SDID + state \times year FE specification is estimated on the 39 mixed states (those with both treated and control counties), retaining 85% of treated and 92% of control counties; the other three SDID specifications use the full balanced 2003–2011 sample.

Appendix A: Alternative specifications and additional results

This appendix presents a series of sensitivity analyses for the difference-in-differences estimates presented in Figure 3 (entropy-balanced Poisson) and Figure 4 (SDID): the post-gestation ATT across five estimator–weighting combinations and four nested control specifications, plus a continuous-coverage Poisson event study on the full county universe (Figure A.3, Figure A.4); the SDID estimates in levels rather than logs (Figure A.5); alternative treatment/control coverage cutoffs (Figure A.6); dropping counties with T-Mobile service to purge the only pre-2009 non-iPhone smartphone channel from both arms (Figure A.7); re-estimating the SDID on a narrow 2003–2009 window that pre-dates the late-2009 arrival of Android on Sprint and Verizon (Figure A.8); and the entropy-balanced Poisson Specification 4 counterpart to the subgroup forest plot in the body (Figure A.9).

A.1 Robustness across estimators and control sets

Figure A.1 and Figure A.2 present the full event-study evolution of the entropy-balanced Poisson and SDID results, respectively, under each of the four nested specifications of controls described in the main text: (1) no time-varying covariates; (2) Black, Hispanic, and other-race shares interacted with year; (3) + county-level economic conditions and state-level reproductive- and welfare-policy exposures; and (4) + state-by-year fixed effects. The estimated effects are consistent across all four specifications for both the entropy-balanced Poisson (Figure A.1) and SDID (Figure A.2).

Figure A.3 arrays the post-gestation ATT (mean of the 2008–2011 event-time coefficients on the log birth rate) across five estimator–weighting combinations—SDID, entropy-balanced Poisson, unweighted Poisson, entropy-balanced OLS, and pop-weighted OLS—and, within each, across the four specifications. The first two estimator rows—SDID and entropy-balanced Poisson—summarize the per-specification post-gestation ATTs visualized in Figure A.1 and Figure A.2. As noted in the body, the SDID estimates are typically larger in magnitude and more precise than those from the other estimators, but the pattern is robust across the 20 cells per age. The estimated ATT for 15–19 year-olds ranges from -4.1% to -8.7% , and that for 20–24 year-olds from -3.1% to -6.6% , with every cell negative and statistically distinguishable from zero at the 5% level.

Figure A.4 reports a complementary check that relaxes the binary bucketing entirely—a Poisson event study in the continuous AT&T coverage measure on the full county universe. These results also are consistent with those of the binary models. For instance, the average post-treatment ATT for 15–19 year-olds is -0.047 (p?).

A.2 SDID estimates of birth rate outcomes in levels

Figure A.5 reproduces the SDID event study with the birth rate in levels (births per 1,000 women) rather than logs. The level estimates parallel the log results in Figure 4 at the ages where the SDID finds substantial effects: post-gestation ATTs of -2.63 ($p < 0.01$) births per 1,000 women at ages 15–19, -5.71 ($p < 0.01$) at 20–24, -3.31 ($p < 0.01$) at 30–34, and -0.82 ($p < 0.01$) at 35–39, which against the 2007 baseline birth rates in treated counties (42.1, 121.2, 97.2, and 41.9 per 1,000) imply proportional declines of roughly 6%, 5%, 3%, and 2%—in line with the log point estimates of -0.084 , -0.069 , -0.036 , and -0.014 . The 25–29 and 40–44 estimates are small and statistically indistinguishable from zero in levels (-0.88 and $+0.21$ per 1,000), consistent with the diffuse pattern at these ages already flagged in the body.

A.3 Sensitivity to the coverage cutoff

The main specification buckets counties as treated ($>90\%$ AT&T coverage) versus control ($<10\%$). Figure A.6 re-runs the SDID event study at a looser cutoff ($>80\%/<20\%$ AT&T) and a tighter cutoff ($>95\%/<05\%$ AT&T) and overlays all three event-study paths in each age panel, so the full dynamic response can be inspected rather than just the post-gestation summary. [TO DO: note stability of magnitude and the expected widening of CIs at the tight cutoff.]

A.4 Dropping counties with T-Mobile service

A concern for the late post-period is that control counties (those with low AT&T coverage) could obtain smartphones through a rival carrier rather than through the iPhone, attenuating the treated–control contrast that the SDID estimates. The first Android phone, the HTC Dream / T-Mobile G1, launched on October 22, 2008 and was T-Mobile-exclusive; the Sprint HTC Hero arrived on October 11, 2009 and the Verizon Motorola Droid on November 6, 2009. AT&T was last to Android, with

the Motorola Backflip on March 7, 2010, reflecting its iPhone exclusive. Counties without T-Mobile service therefore had no Android channel of any kind before late 2009. As a robustness check we re-estimate the SDID event study after additionally restricting both arms to counties with less than 10% T-Mobile 3G coverage, eliminating the only pre-2009 non-iPhone smartphone option from the sample (Figure A.7).

The restriction prunes the sample asymmetrically. Control counties (AT&T coverage <10%) were already nearly T-Mobile-free—98% of them have T-Mobile coverage below 10%—so the additional cut removes only 24 of the 1,399 controls. Treated counties (AT&T >90%), in contrast, skew urban and tend to have T-Mobile as well: only 40% have T-Mobile coverage below 10%, so the cut reduces the treated arm from 914 counties to 369. The check is thus less a contamination scrub of the controls than a stress test that asks whether the main result holds on the more rural subset of treated counties that never had any non-iPhone Android option.

The teen and young-adult ATTs survive cleanly. The post-gestation ATT is -0.041 ($p < 0.01$) at ages 15–19 and -0.038 ($p < 0.01$) at ages 20–24, statistically and substantively similar to the corresponding main estimates, indicating that the iPhone-era fertility decline among teens and young adults is not an artifact of early-Android substitution in the control group. The 25–29 effect, however, attenuates to -0.004 (0.006), no longer statistically distinguishable from zero. We read this as evidence that the 25–29 channel in the main specification ran disproportionately through the urban high-T-Mobile counties that this restriction drops—a compositional rather than a contamination story—and treat the 25–29 effect with corresponding caution. At ages 30 and above the no-T-Mobile estimates are mixed (a clean -0.025 ($p < 0.01$) at 30–34, otherwise small and imprecisely estimated), in the same diffuse pattern flagged by the Verizon placebo (Figure B.4), reinforcing our decision to anchor the paper on the 15–24 results.

A.5 Narrow 2003–2009 window: pre-dating the Sprint and Verizon Android launches

The robustness check in Figure A.7 purges T-Mobile counties to eliminate the pre-2009 G1 Android channel from both arms. A complementary, more direct address of the rival-carrier-Android concern is to truncate the post-period to birth years conceived before Android arrived on Sprint and Verizon. The Sprint HTC Hero launched on October 11, 2009 and the Verizon Motorola Droid on November 6,

2009; the conceptions affected by those launches appear in 2010 and later birth years. We therefore re-estimate the SDID on a balanced 2003–2009 panel that excludes both 2010 and 2011, leaving 2008–2009 as the post-period (Figure A.8). All other features—the >90%/<10% AT&T sample, the 2008 first treated year, the SDID estimator, and the county-clustered bootstrap with 500 replications—match Figure 4. The post-gestation ATT is now the mean of the 2008 and 2009 event-time coefficients on the log birth rate.

Figure A.8 reports the narrow-window post-gestation ATTs across the six age groups. At the headline ages, the estimates remain negative and statistically significant: -0.060 ($p < 0.01$) at 15–19 and -0.043 ($p < 0.01$) at 20–24. Both are smaller in magnitude than the full-window headline (-0.084 and -0.068), consistent with iPhone adoption compounding through 2010–2011 as ownership in treated counties continued to rise. The older-age narrow estimates are similar in magnitude to the headline (-0.008 at 25–29, -0.027 ($p < 0.01$) at 30–34, -0.019 ($p = 0.03$) at 35–39, and -0.037 ($p = 0.02$) at 40–44). The pattern as a whole—negative coefficients at every age, with sign and rough magnitude preserved relative to the headline—indicates that the main result is not driven by late-period Android substitution in control counties.

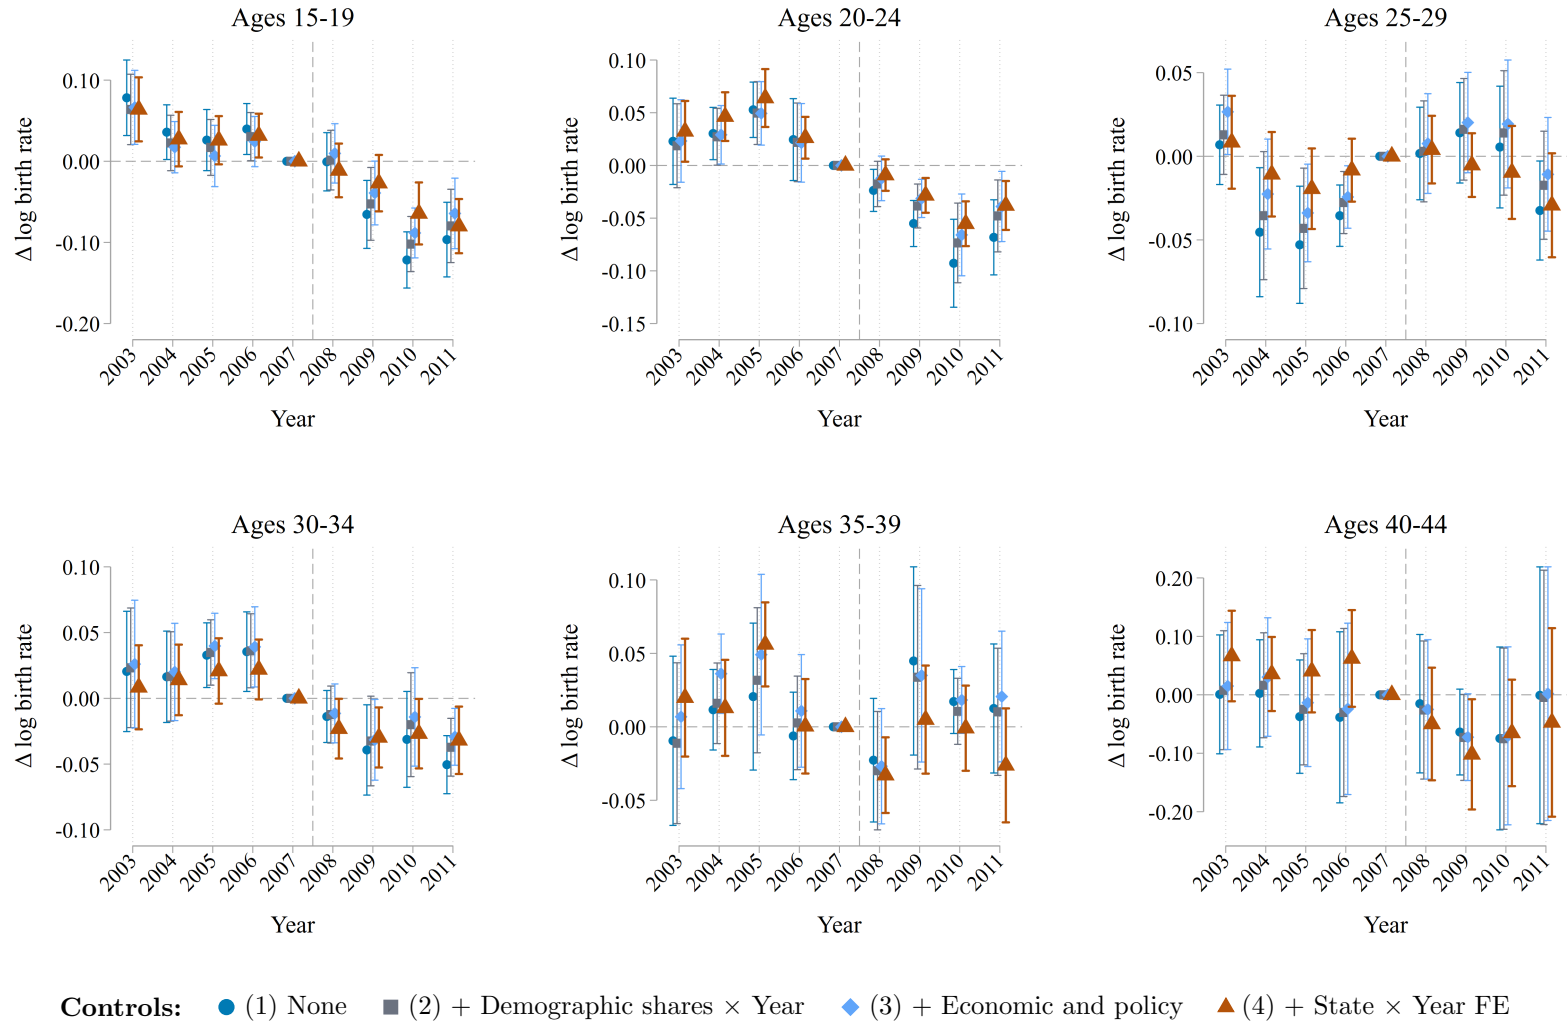
A.6 Subgroup post-gestation ATT: entropy-balanced Poisson Specification 4

Figure A.9 reproduces the subgroup forest plot in Figure 5 using the entropy-balanced Poisson Specification 4 estimator (demographic shares interacted with year, county-level economic controls, and state-by-year fixed effects) rather than the SDID. For the six age rows, the entropy-balanced Poisson uses the per-age entropy-balance weights from Figure 3, so those rows reproduce the headline entropy-balanced Poisson Specification 4 estimates. For the twelve non-age rows, a single women-aged-15–44 entropy-balance weight is computed once and applied uniformly across the twelve regressions; balancing covariates (Black, Hispanic, urban, and Republican shares) match the headline entropy-balanced Poisson, with the six urban×broad-region strata.

As observed with other comparisons of entropy-balanced Poisson and SDID estimates, these too are less precise. But the general pattern of heterogeneous effects remain evident, with larger point estimates of the ATT for younger populations. The exception is a large but highly imprecise estimated effect for women aged 40–44

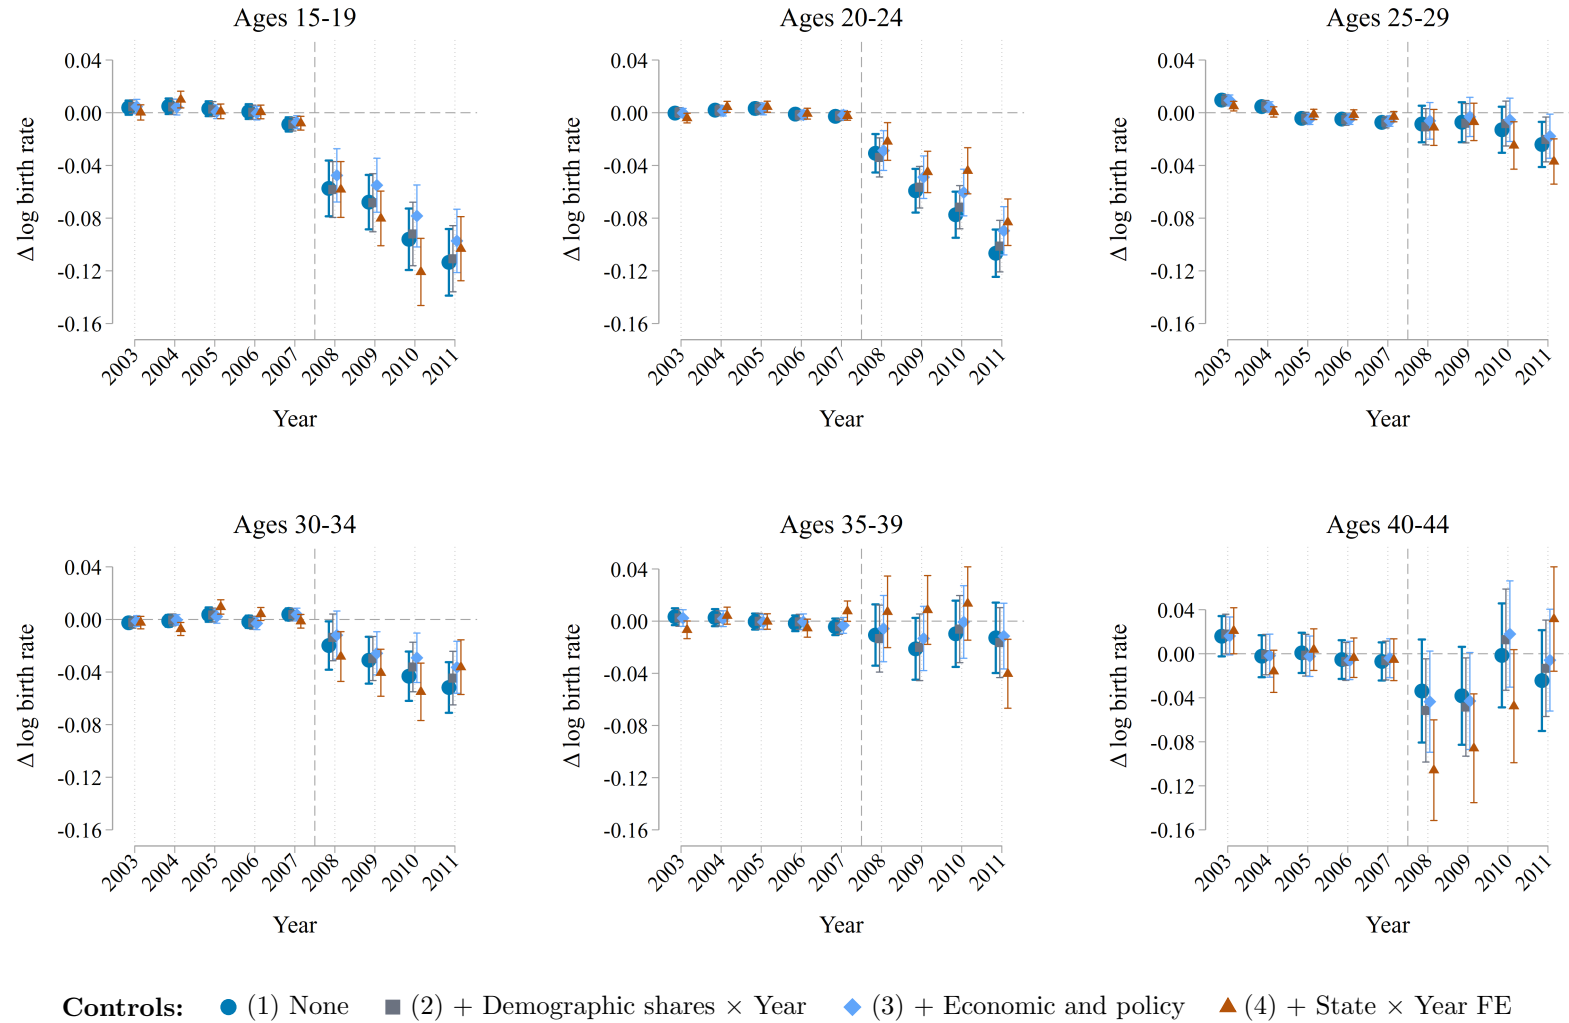
of -6.4% ($p = 0.16$), with a 95% confidence interval of $[-14.6\%, +2.6\%]$.

Figure A.1: Estimated effect of AT&T 3G coverage on birth rates, by age, from an entropy-balanced Poisson event study under four nested control specifications




Notes: Companion to Figure 3; each panel overlays the entropy-balanced Poisson event study under the four nested control specifications shown in the legend, with horizontal offsets within each year that are graphical only. Specification 4, which adds state \times year fixed effects, is the headline specification reported in Figure 3. Post-gestation ATTs for all four specifications are tabulated in Table 2. All four specifications use the full entropy-balanced sample. Standard errors are clustered at the county level. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure A.2: Estimated effect of AT&T 3G coverage on birth rates, by age, from a synthetic difference-in-differences event study under four nested control specifications



Notes: Companion to Figure 4; each panel overlays the SDID event study under the four nested control specifications shown in the legend, with horizontal offsets within each year that are graphical only. Specification 1, which uses no covariates, is the headline specification reported in Figure 4. Post-gestation ATTs for all four specifications are tabulated in Table 2. Specification 4 adds state \times year fixed effects, which require within-state treated-versus-control variation and therefore restrict the sample to the 39 “mixed” states with both treated and control counties (85% of treated counties, 92% of control counties); the other three specifications use the full balanced sample. Inference is by county-clustered bootstrap with 500 replications. Sources: NCHS (2025), SEER (2026), NTIA (2010).


Figure A.3: Estimated effect of AT&T 3G coverage on birth rates across estimators and control specifications, by age



Controls: ● (1) None ■ (2) + Demographic shares × Year ◆ (3) + Economic and policy ▲ (4) + State × Year FE

(continued on next page)


Figure A.3: (continued)



Controls: ● (1) None ■ (2) + Demographic shares × Year ◆ (3) + Economic and policy ▲ (4) + State × Year FE

(continued on next page)

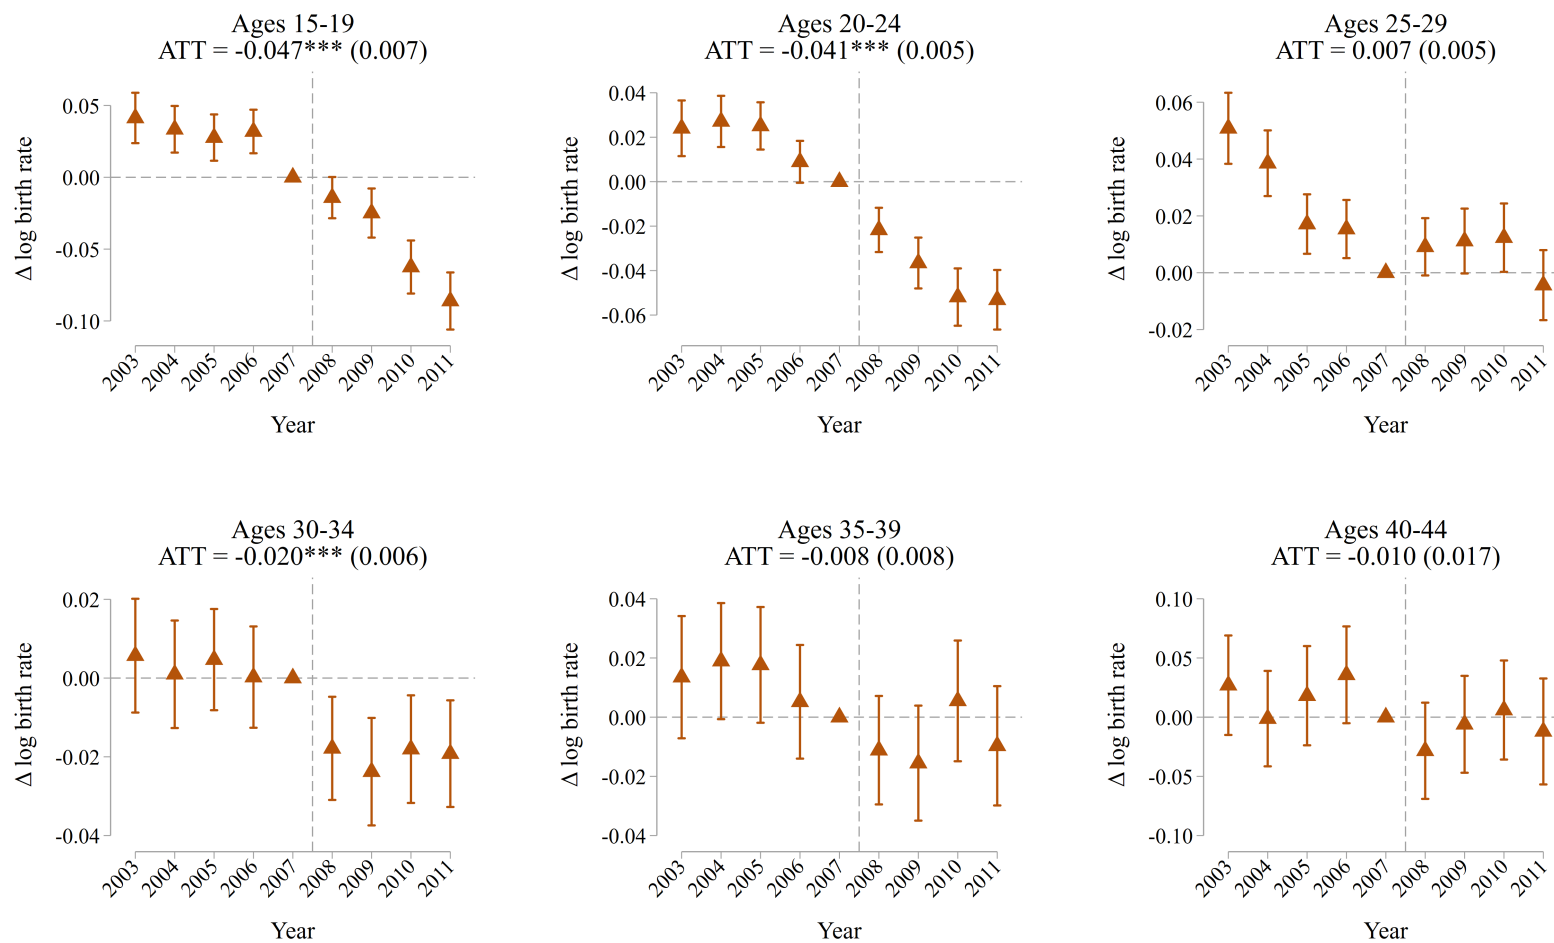
Figure A.3: (continued)



Controls: ● (1) None ■ (2) + Demographic shares \times Year ◆ (3) + Economic and policy ▲ (4) + State \times Year FE

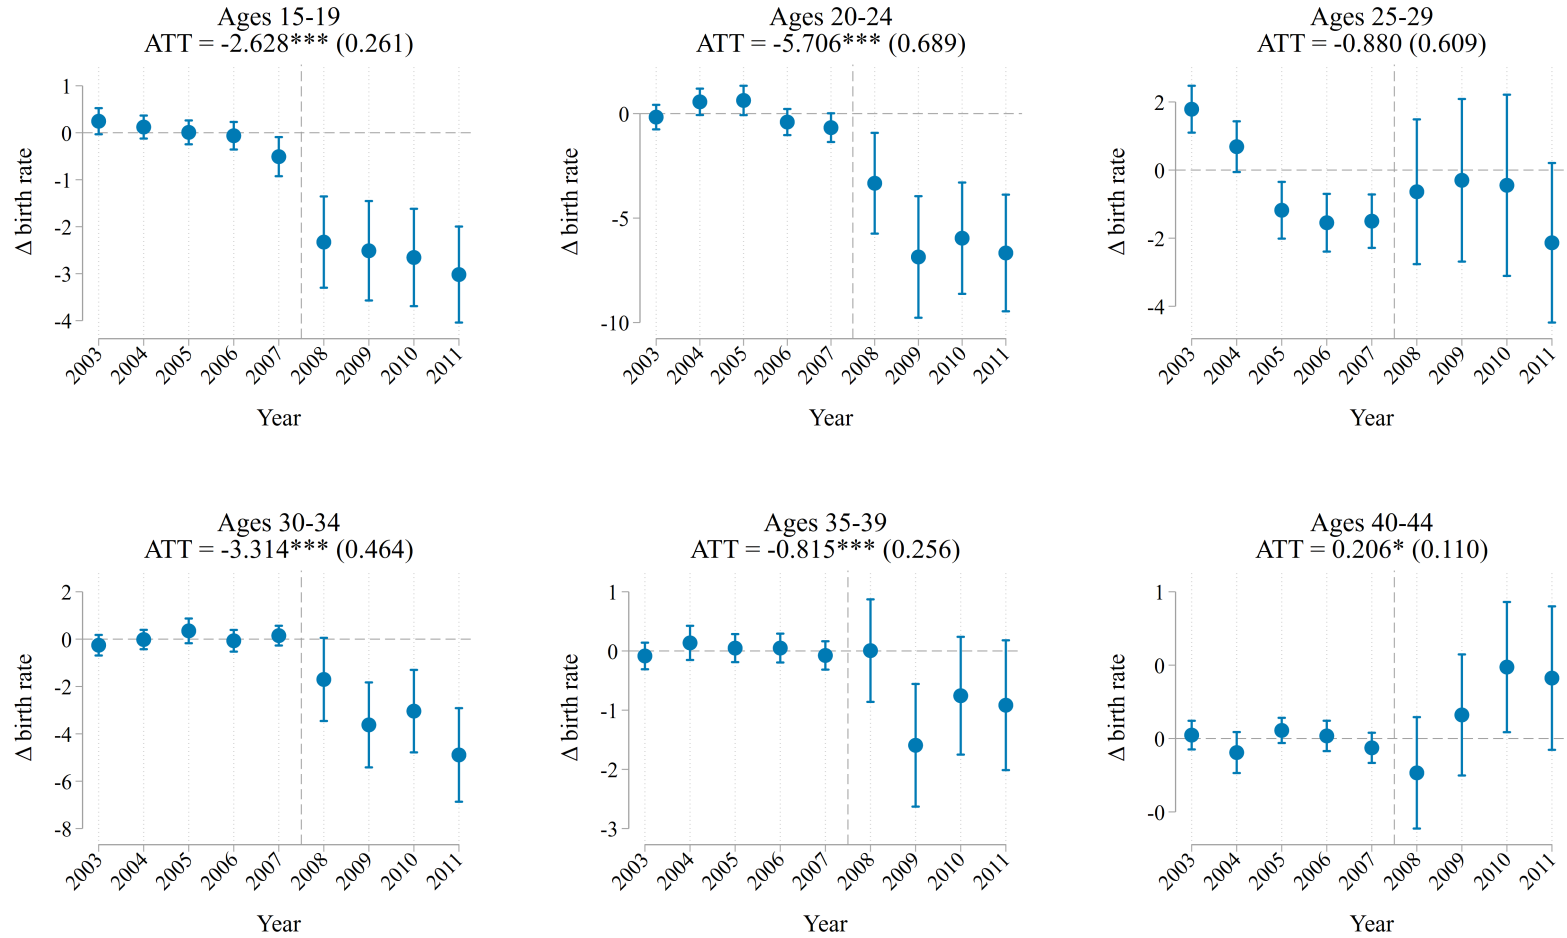
Notes: Each panel plots the post-gestation ATT (mean of the 2008–2011 event-time coefficients on the log birth rate) with 95% confidence intervals, across five estimators and four nested control specifications. SDID denotes synthetic difference-in-differences (Specification 1, no covariates, is reported in Figure 4); Poisson denotes the entropy-balanced PPML estimator (Specification 4, with state \times year fixed effects, is reported in Figure 3); Unweighted Poisson is the corresponding PPML without entropy weights; Entropy-balanced OLS uses population aweights with the per-age entropy-balance weights of Figure 3 layered on top; Pop-weighted OLS uses population aweights only, with no entropy balancing. SDID Specification 4 is estimated on the 39 mixed states with both treated and control counties (see Figure A.2 notes); all other rows use the full balanced sample. All models share the binary $>90\%/<10\%$ AT&T treated/control sample. Standard errors are clustered at the county level; SDID inference uses a 500-replication bootstrap. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure A.4: Estimated effect of continuous AT&T 3G coverage on birth rates, by age, from a Poisson event study



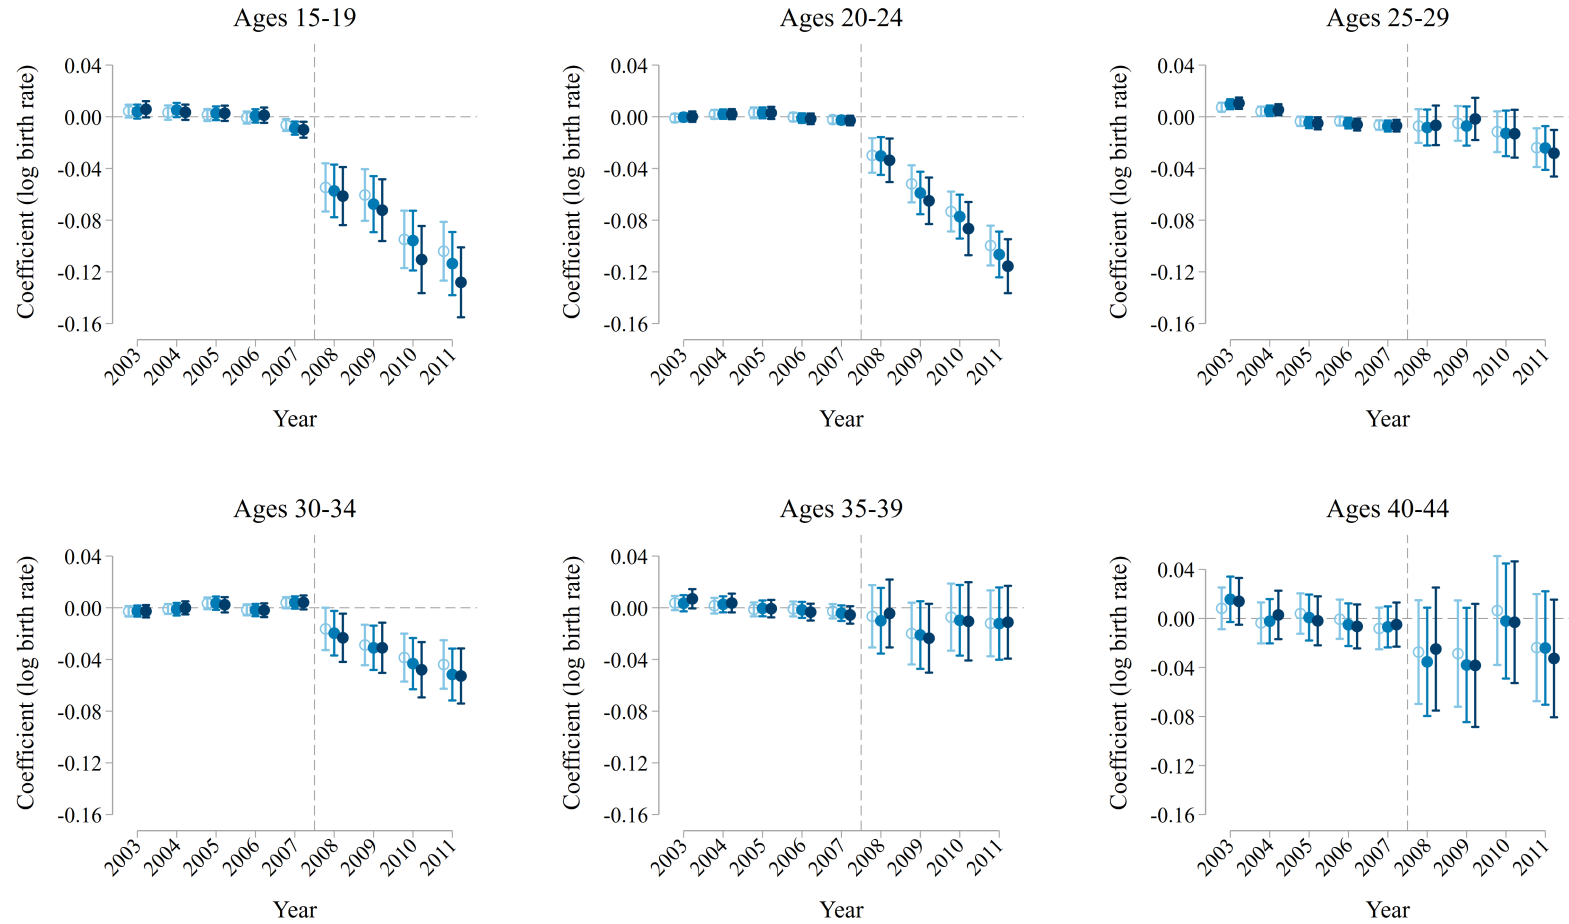
Notes: Companion to Figure 3; each panel reports a Poisson (PPML) event study of age-specific births (exposure: population) on the *continuous* county AT&T coverage share interacted with birth year, 2007 reference. The specification matches Specification 4 of Figure A.1—demographic shares \times year, economic controls, and state \times year fixed effects (which absorb the state-level policy controls)—and is the headline specification reported in Figure 3. The sample is the full county universe, without the $>90\%/<10\%$ AT&T bucketing applied in Figure 3. Standard errors are clustered at the county level; bars are 95% confidence intervals. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure A.5: Estimated effect of AT&T 3G coverage on birth rates in levels, by age, from a synthetic difference-in-differences event study



Notes: As Figure 4, but the outcome is the age-specific birth rate in levels (births per 1,000 women) rather than its log. Bars are 95% confidence intervals from a county-clustered bootstrap with 500 replications. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure A.6: Estimated effect of AT&T 3G coverage on birth rates, by age, from a synthetic difference-in-differences event study under three treatment/control cutoffs

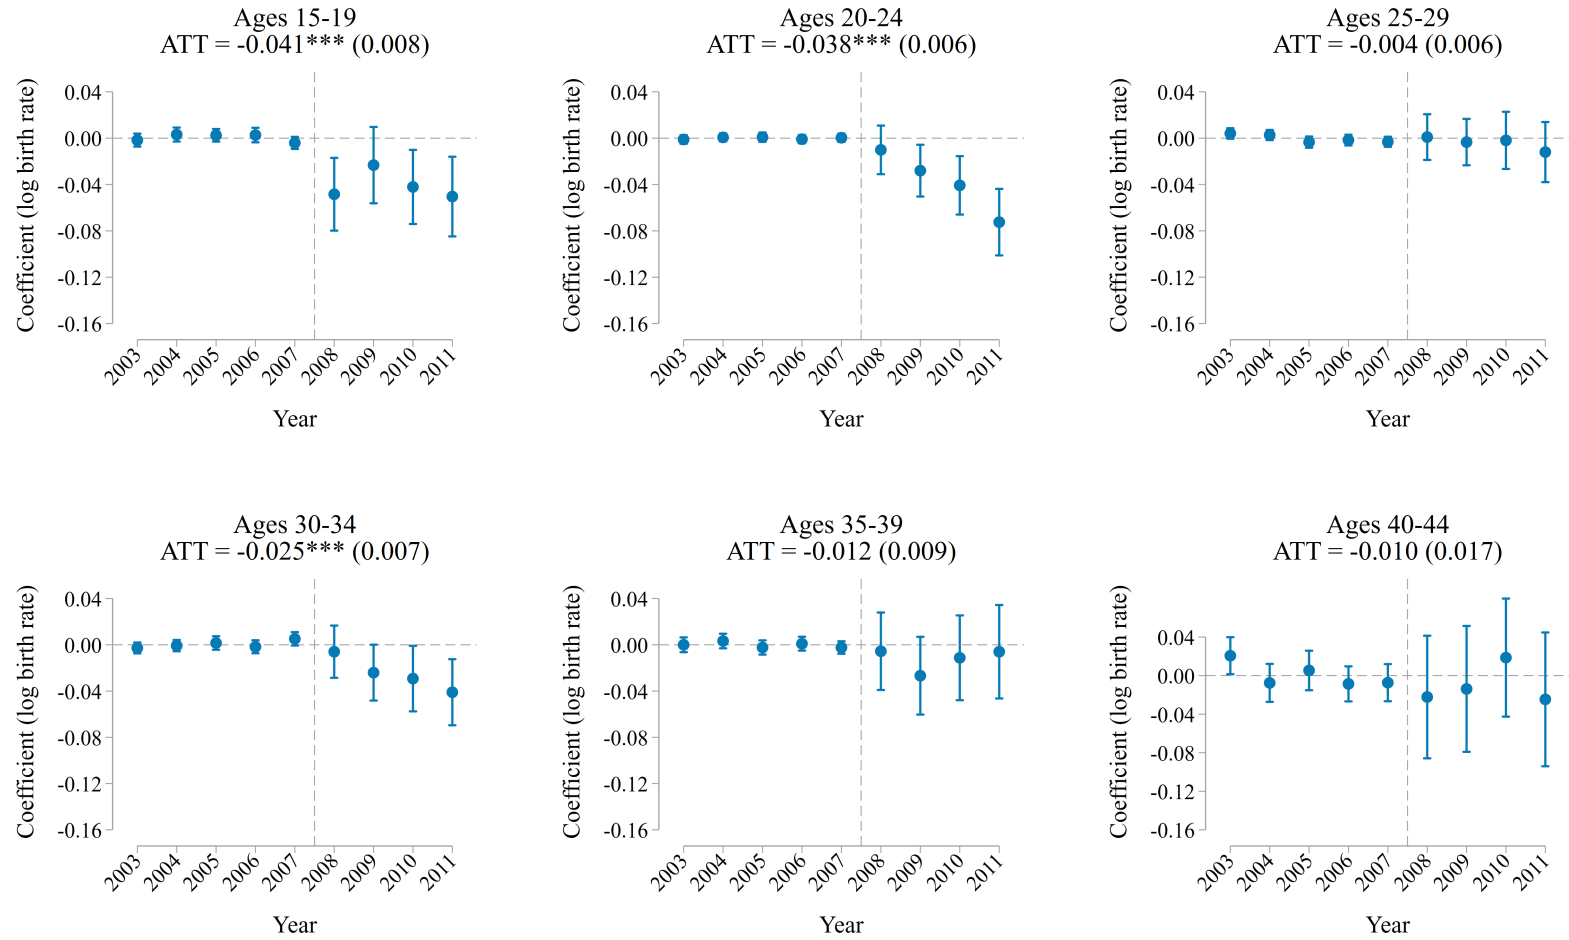


Three alternative thresholds to define treatment and control groups:

- Loose (>80% / <20%)
- Main (>90% / <10%)
- Tight (>95% / <05%)

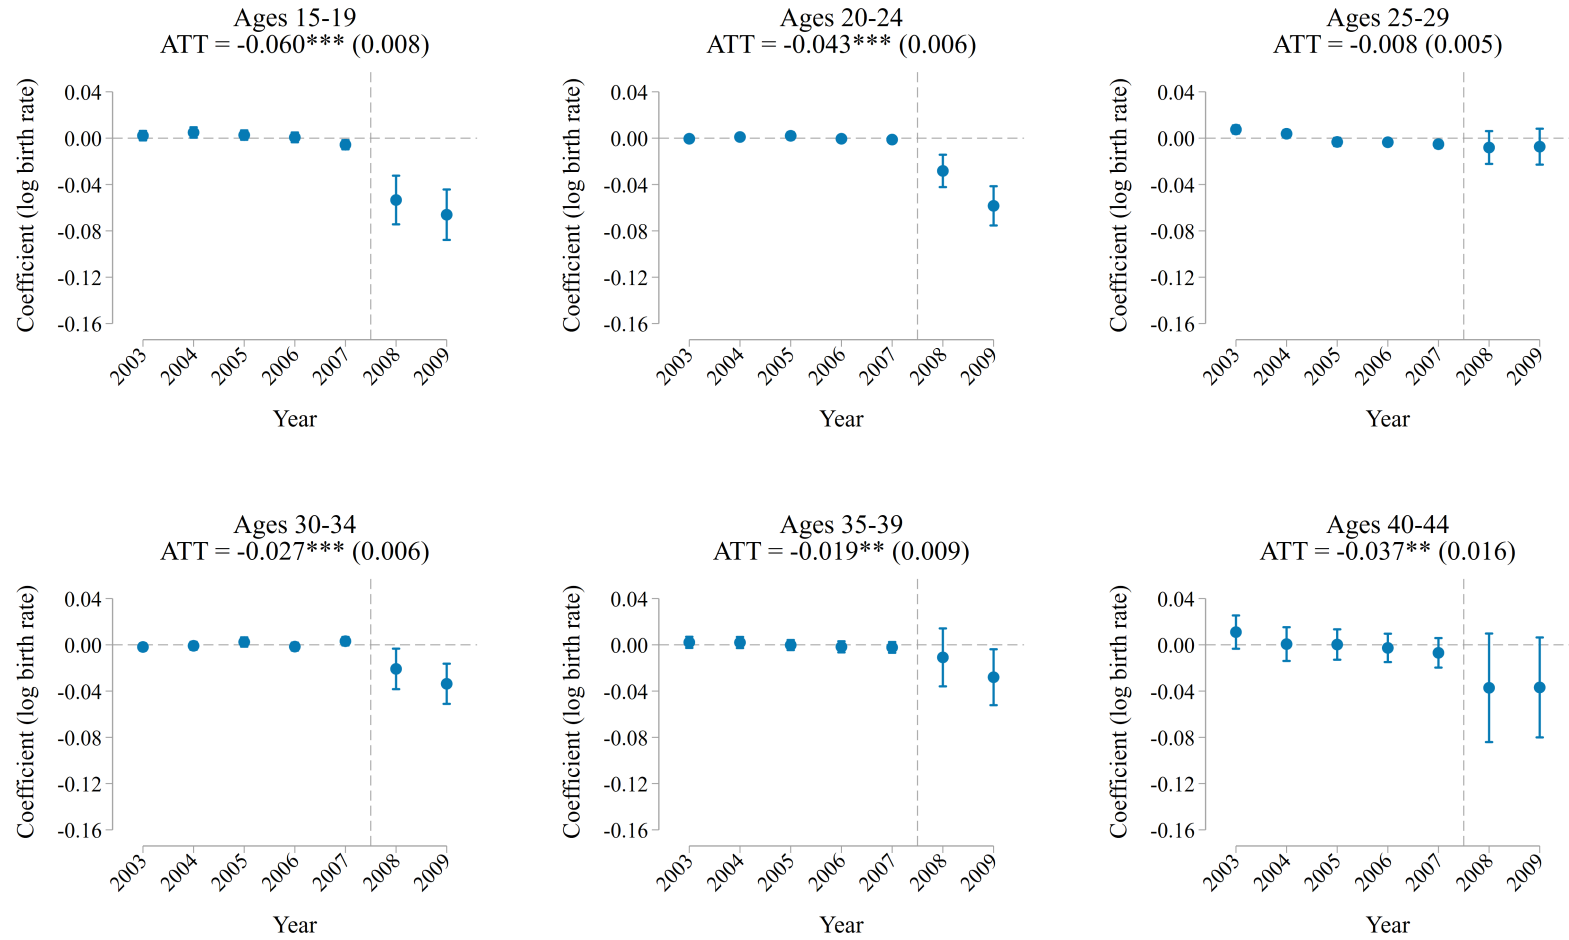
Notes: Each panel overlays the SDID event study (Figure 4) at the three treatment/control cutoffs shown in the legend, with horizontal offsets within each year that are graphical only. As the cut loosens, more marginal counties enter and the treatment contrast weakens; as it tightens the contrast sharpens but N falls and CIs widen. All other features—outcome (log birth rate), 2003–2011 window, 2008 first treated year, county-clustered bootstrap with 500 replications, post-gestation ATT averaged over 2008–2011—match Figure 4. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure A.7: Estimated effect of AT&T 3G coverage on birth rates, by age, from a synthetic difference-in-differences event study dropping counties with T-Mobile service



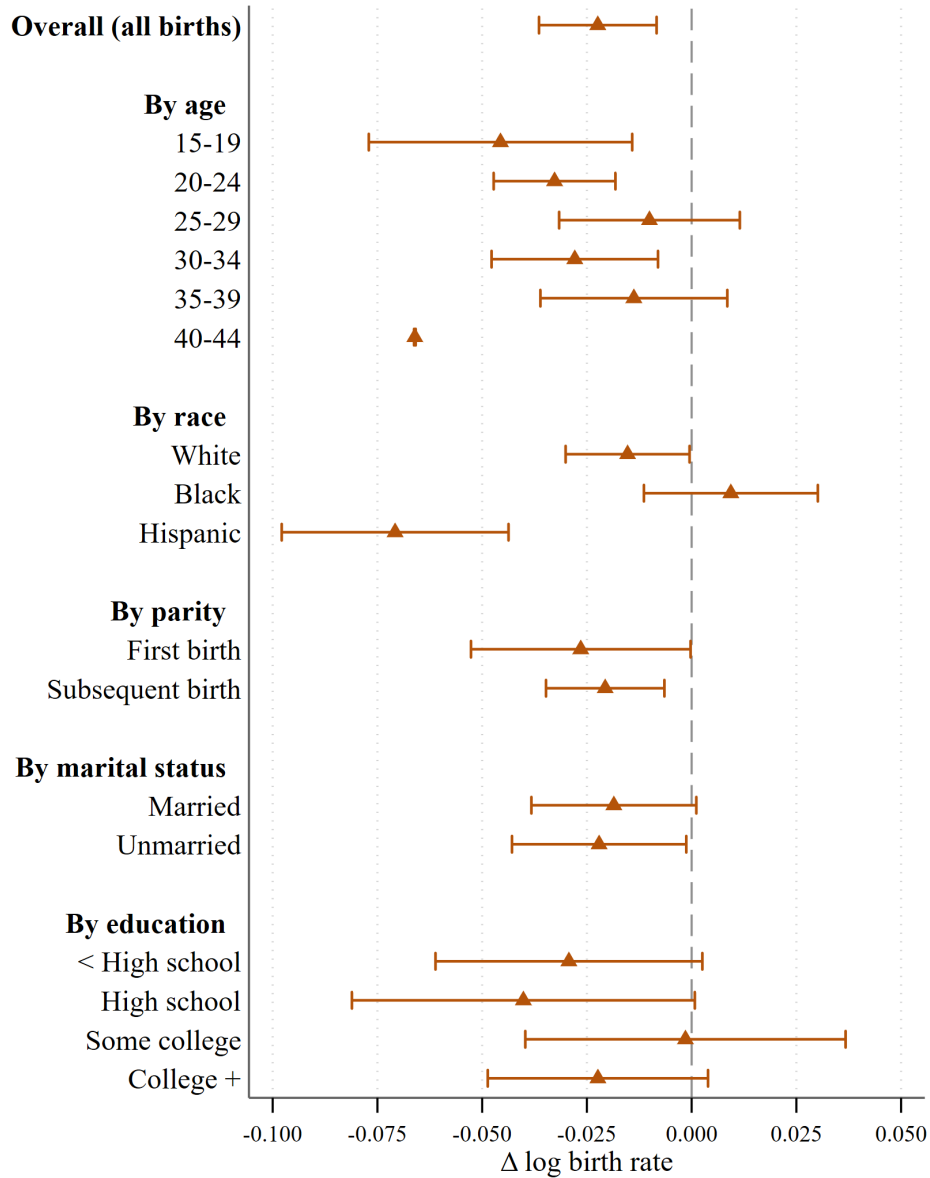
Notes: As Figure 4, but the sample is additionally restricted to counties with less than 10% T-Mobile 3G coverage. The T-Mobile G1 (October 2008) was the first Android phone and the only Android channel before late 2009, when the Sprint HTC Hero (October 2009) and Verizon Motorola Droid (November 2009) launched; dropping T-Mobile counties therefore purges the only pre-2009 non-iPhone smartphone substitute from both arms. Treated counties drop from 914 to 369 (the restriction selects against urban counties); control counties from 1,399 to 1,375 (low-AT&T counties were already nearly T-Mobile-free). Bars are 95% confidence intervals from a county-clustered bootstrap with 500 replications. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure A.8: Estimated effect of AT&T 3G coverage on birth rates, by age, from a synthetic difference-in-differences event study on a 2003–2009 window



Notes: As Figure 4, but the balanced county panel is truncated to 2003–2009 so the post-period (2008–2009) excludes birth years materially exposed to the Sprint HTC Hero (Oct 2009) and Verizon Motorola Droid (Nov 2009) Android launches. The post-gestation ATT reported in each subplot title is the mean of the 2008 and 2009 event-time coefficients on the log birth rate. Bars are 95% confidence intervals from a county-clustered bootstrap with 500 replications. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure A.9: Estimated effect of AT&T 3G coverage on birth rates, by subgroup, from an entropy-balanced Poisson event study



Notes: Each marker is the post-gestation mean of the 2008–2011 event-time coefficients from the entropy-balanced Poisson Specification 4 estimator on the log birth rate, with delta-method 95% confidence intervals (county-clustered). The six age rows reproduce the by-age estimates in Figure 3 and use its per-age entropy-balance weights; the twelve non-age rows use a single shared women-aged-15–44 entropy balance. Age- and race-specific rows use group-specific population denominators; the overall, parity, marital-status, and education rows use women aged 15–44 as the denominator. Because NCHS removed the unrevised-certificate maternal-education item beginning in 2011—leaving education unreported for births in states still using the 1989 certificate—the education rows are estimated on the 2003–2010 window; all other rows use the full 2003–2011 window. The 95% confidence interval for ages 40–44 (true CI -0.157 to $+0.025$) extends beyond the chart range and is suppressed; only the point estimate is shown. A parallel SDID version of this figure is reported in Figure 5. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Appendix B: Identification diagnostics for the synthetic difference-in-differences estimates

We perform two kinds of placebo test on the research design. [Section 9.1](#) reports placebo-in-time tests that artificially reassign the treatment date to years before the iPhone launch. [Section 9.2](#) reports cross-carrier placebos that apply the design to counties served by Verizon ([Figure B.4](#)) and by Sprint ([Figure B.6](#)). The first probes whether the SDID estimator manufactures effects from pre-2007 dynamics; the second probes whether the design picks up a generic 3G-broadband effect rather than something iPhone-specific.

B.1 Placebo-in-time tests

[Figure B.1](#) presents a placebo-in-time test that re-runs the SDID estimator with the treatment date reassigned to 2003, 2004, 2005, or 2006. A genuine effect at 2008 should leave the ATT for these falsified dates near zero. Each panel plots the first post-period ($t = 1$) coefficient at each falsified introduction year (2003–2006), estimated on the pre-iPhone sample, alongside the true 2008 first-year coefficient we estimate in the paper. The placebo coefficients are smaller and the true 2008 estimate is visibly separated from the placebo cloud at the young ages where the headline effect is largest.

[Figure B.2](#) reports the entropy-balanced Poisson Specification 4 placebo. Several placebo coefficients are sizable: at ages 15–19 the 2004 placebo is -3.7% ($p = 0.05$); at ages 20–24 the 2006 placebo is -3.9% ($p < 0.01$); at ages 35–39 the 2006 placebo is -5.4% ($p < 0.01$). At five of six age bands the largest placebo magnitude exceeds the true 2008 first-year coefficient, leaving the headline year-by-year signal not clearly distinguishable from the placebo distribution.

The contrast between the two figures is itself a piece of evidence for preferring the SDID as our central estimator: its data-driven control-county and pre-period weights yield a signal that separates from the year-by-year placebo noise, whereas the entropy-balanced Poisson’s parametric controls and fixed entropy-balance weights only neutralize the pre-period heterogeneity at the four-year-mean level reported in [Figure 3](#).

B.2 Carrier placebos

Figure B.4 and Figure B.6 apply the SDID design to counties served by Verizon and by Sprint, holding fixed the 2008 first-treated-year coding used for AT&T. If the design were capturing something general about 3G-broadband rollout rather than the iPhone specifically, these cross-carrier comparisons should themselves produce “effects” during the AT&T-exclusivity window. The two placebos require careful interpretation, because both carriers launched Android smartphones within our 2003–2011 birth-year window: the Verizon Motorola Droid on November 6, 2009 and the Sprint HTC Hero on October 11, 2009. Given the ~ 9 -month gestation lag, the first birth year materially exposed to either rival-carrier Android is 2010, and 2011 is the only fully exposed year. We therefore split each cross-carrier placebo into two event-window estimates. The “2008–2009 ATT (placebo)” averages the 2008 and 2009 event-time coefficients—the fully treated pre-Android window. This *is* a placebo: under the iPhone-specific reading of our main result, a 3G-coverage advantage on a carrier that lacks both the iPhone and Android should produce no effect on age-specific birth rates in this window, and a null estimate validates the design. The “2011 ATT (Android)” is the 2011 coefficient alone—the only birth year fully exposed to the rival-carrier Android. This *is not* a placebo: it estimates the causal effect of rival-carrier Android adoption on age-specific birth rates over a single year, using the same SDID comparison machinery as the main AT&T estimate. A negative 2011 ATT at the ages where the main AT&T effect is largest, far from threatening the iPhone interpretation, would corroborate the broader reading that the relevant channel is *smartphone access*, with the iPhone (on AT&T) and Android (on Verizon and Sprint) providing independent vehicles for the same underlying phenomenon.

Figure B.4 reports the Verizon placebo, with the underlying county classification mapped in Figure B.3. We assign the SDID “treatment” to counties with $>90\%$ Verizon 3G coverage²³ and $<10\%$ AT&T coverage, using counties with $<10\%$ cover-

²³Verizon Wireless completed its acquisition of Alltel in January 2009, but in the December 2010 SBI snapshot Alltel still reports under its own holding-company name (HOCONUM 900046 / HOCONAME “Alltel Wireless”). We aggregate both “Verizon Communications Inc.” and “Alltel Wireless” under our Verizon classifier to capture the consolidated Verizon footprint at the snapshot date. DOJ-mandated divestitures from the Verizon–Alltel merger went to Atlantic Tele-Network (ATN), which operated those rural markets under the Alltel brand through 2013; the SBI does not contaminate our aggregation, since all 240 counties with “Alltel Wireless” filings in December 2010 are also counties where “Verizon Communications Inc.” files natively, with no county in the Alltel-only pattern an ATN divested-rural footprint would produce.

age from both carriers as controls; AT&T-served counties are excluded so the AT&T effect cannot leak in. At ages 15–19 and 20–24, where the main AT&T effects are largest, the 2008–2009 placebo is precisely null (+0.006 and -0.004 ; neither significant), supporting an iPhone-specific reading during the AT&T-exclusivity window. The 2011 Android estimate is negative at 15–19 (-0.033 , $t \approx 1.4$) and null at 20–24 (+0.013); the teen 2011 point estimate is consistent with the gestation-adjusted arrival of Verizon Android, though it falls short of conventional significance. At ages 25–44 the placebo period itself produces statistically significant negative coefficients at 25–39 (-0.029 to -0.048) and a negative but imprecisely estimated coefficient at 40–44 (-0.035 , n.s.), indicating that connectivity-tier differences among rural counties are associated with diverging older-age fertility for reasons unrelated to the iPhone. We therefore interpret the older-age AT&T estimates with caution and rest our central claims on the 15–24 results.


Figure B.6 reports an analogous placebo for Sprint, which acquired Android via the HTC Hero on October 11, 2009 and the iPhone on October 14, 2011 (after every birth-year conception window in our sample). The Sprint county classification is mapped in Figure B.5.²⁴ At ages 15–19 the Sprint 2008–2009 placebo is null (-0.014 , n.s.) and the 2011 Android estimate is even more sharply negative than Verizon’s: -0.080 ($p < 0.01$). At ages 20–24 the placebo is statistically significant (-0.031 , $p = 0.01$) and the 2011 Android estimate is -0.080 ($p < 0.01$), indicating a Sprint-Android effect at this age that the Verizon placebo did not pick up. At 25–29 both estimates are exactly null. The older-age pattern is noisier and qualitatively similar to the Verizon placebo, reinforcing our older-age caution.

Taken together, the placebo tests support our main reading. The SDID placebo-in-time confirms the estimator is not manufacturing an effect from pre-2007 dynamics. The 2008–2009 cross-carrier placebos at 15–19 are precisely null on both Verizon and Sprint, confirming that the design does not pick up a generic 3G-coverage effect during the AT&T-exclusivity window. The 2011 Android estimates—particularly the gestation-aligned Sprint-Hero teen effect (-0.080 , $p < 0.01$), with the Verizon-Droid teen effect in the same direction but imprecisely estimated (-0.033)—reinforce rather than threaten the iPhone interpretation, documenting the same underlying

²⁴We would have preferred T-Mobile, whose G1 (October 22, 2008) was the earliest Android phone, but T-Mobile’s footprint barely overlaps with low-AT&T counties: only a single county satisfies the parallel $>90\%/<10\%$ AT&T cutoff, which itself indicates that the G1 Android channel was structurally absent from the SDID control group. Sprint, by contrast, has 156 counties at the parallel cutoff with a 740-county control arm.


smartphone–fertility channel arriving on rival carriers via Android on exactly the schedule that gestation predicts.

Figure B.1: Estimated effect of AT&T 3G coverage on birth rates at placebo introduction years, by age, from a synthetic difference-in-differences event study




Notes: Each panel reports SDID estimates in which the treatment date is reassigned to a falsified “introduction year” (2003, 2004, 2005, or 2006), estimated on the pre-iPhone sample, alongside the true 2008 estimate (top, in blue). Plotted is the first post-period ($t = 1$) coefficient on the log birth rate with its 95% confidence interval. The sample, weighting, and bootstrap inference match Figure 4. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure B.2: Estimated effect of AT&T 3G coverage on birth rates at placebo introduction years, by age, from an entropy-balanced Poisson event study



Notes: Each panel reports entropy-balanced Poisson Specification 4 estimates in which the treatment date is reassigned to a falsified “introduction year” (2003, 2004, 2005, or 2006), estimated on the pre-iPhone sample with a window mirroring the SDID placebo (four pre-treatment years plus all available post-treatment years through 2007). Plotted is the first post-period ($t = 1$) coefficient on the log birth rate with its 95% confidence interval; the real 2008 row (top, in amber) is from Figure 3. Sample, weighting, and Specification 4 controls match Figure 3; standard errors are clustered at the county level. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Figure B.3: County classification for the Verizon carrier placebo

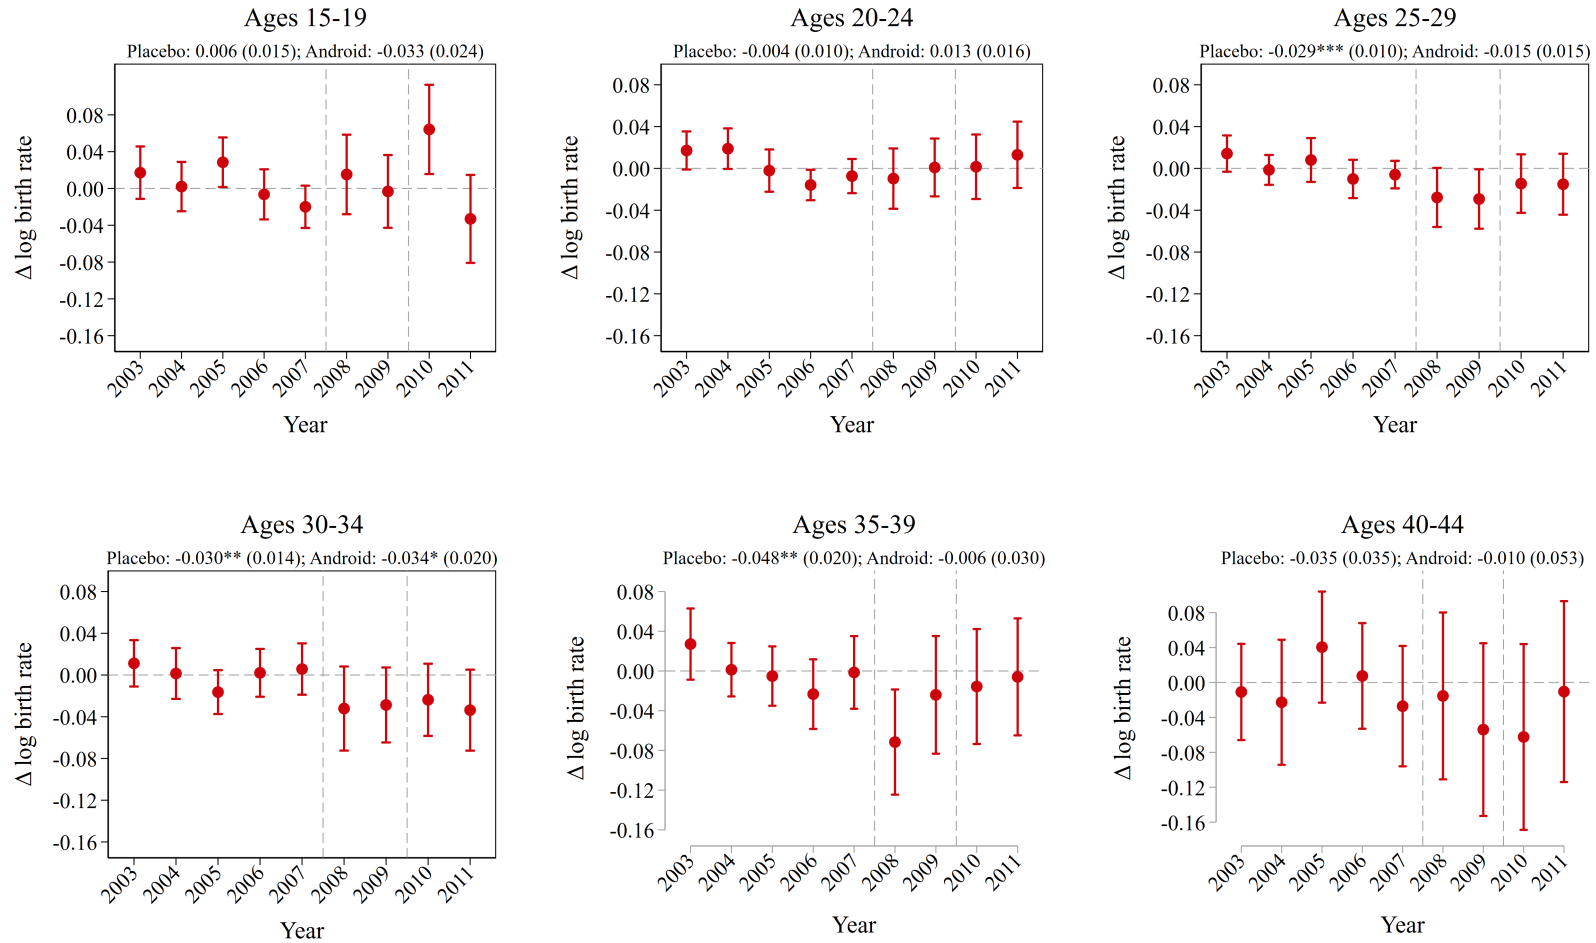


Verizon placebo county classification

□ Excluded □ Control ■ Treated


Notes: County classification for the Verizon carrier placebo whose event-study results appear in [Figure B.4](#). **Treated** counties (red) have >90% Verizon 3G coverage and <10% AT&T coverage. **Control** counties (gray) have <10% coverage from both carriers. **Excluded** counties (white) are everywhere else, including all AT&T-served counties. Alaska and Hawaii are omitted. Source: SBI 2010 (NTIA, 2010).

Figure B.4: Estimated effect of Verizon-only coverage on birth rates, by age, from a synthetic difference-in-differences event study (carrier placebo)



Notes: Each panel reports an SDID event study identical in design to the main result (log age-specific birth rate, balanced 2003–2011 county panel, 2008 first treated year, county-clustered bootstrap with 500 replications), but with a placebo treatment: counties with >90% Verizon 3G coverage and <10% AT&T coverage are coded as “treated” and counties with <10% coverage from both carriers as controls. AT&T-served counties are excluded. The two dashed vertical lines mark the first birth year materially exposed to each carrier’s flagship smartphone given the ~9-month gestation lag: “AT&T iPhone” at the 2006/2007 boundary (the iPhone launched June 29, 2007) and “Verizon Android” at the 2009/2010 boundary (the Motorola Droid launched November 6, 2009). The two summaries above each panel report the mean of the 2008 and 2009 event-time coefficients (“2008–2009 AT&T (placebo)”—the fully treated pre-Android window) and the 2011 coefficient alone (“2011 AT&T (Android)”—the only birth year fully exposed to Verizon Android). The Verizon iPhone arrived February 10, 2011, after the conception window for every birth cohort in the sample, so a Verizon-iPhone effect cannot enter. Markers and confidence intervals are rendered in Verizon red. Bars are 95% confidence intervals from a county-clustered bootstrap with 500 replications. Sources: NCHS (2025) SEER (2026) NTIA (2010)

Figure B.5: County classification for the Sprint carrier placebo

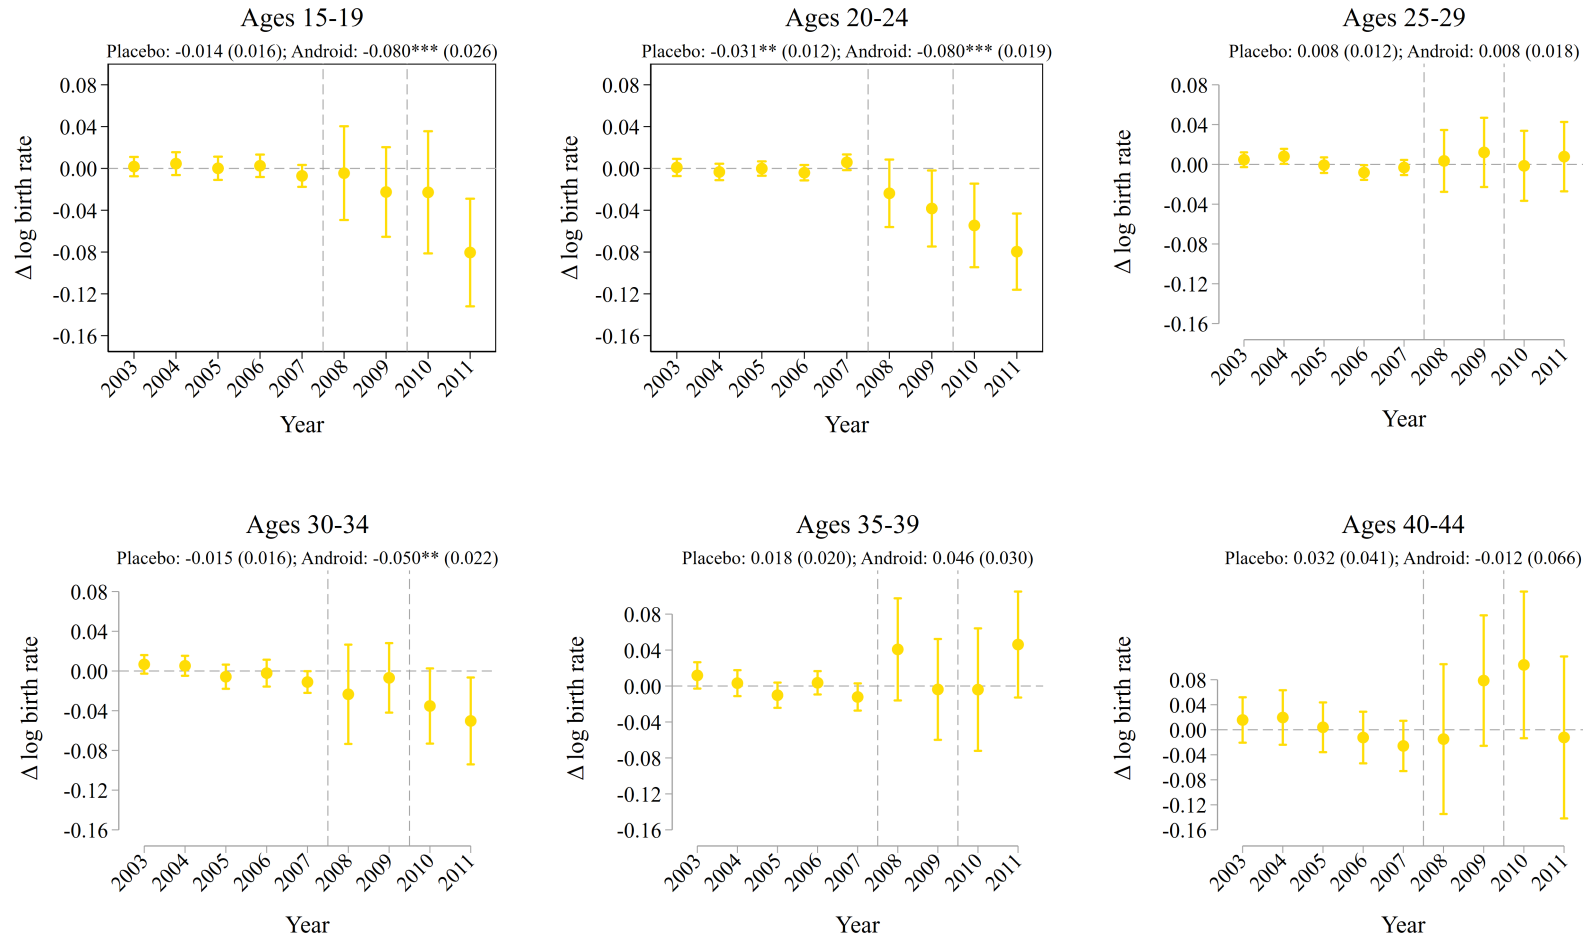


Sprint placebo county classification

□ Excluded □ Control □ Treated

Notes: County classification for the Sprint carrier placebo whose event-study results appear in [Figure B.6](#). **Treated** counties (yellow) have >90% Sprint 3G coverage and <10% AT&T coverage. **Control** counties (gray) have <10% coverage from both carriers. **Excluded** counties (white) are everywhere else, including all AT&T-served counties. Alaska and Hawaii are omitted. Source: SBI 2010 (NTIA, 2010).

Figure B.6: Estimated effect of Sprint-only coverage on birth rates, by age, from a synthetic difference-in-differences event study (carrier placebo)



Notes: Each panel reports an SDID event study identical in design to Figure B.4 except that the placebo carrier is Sprint rather than Verizon: counties with >90% Sprint 3G coverage and <10% AT&T coverage are coded as “treated” and counties with <10% coverage on both as controls. AT&T-served counties are excluded. The two dashed vertical lines mark the first birth year materially exposed to each carrier’s flagship smartphone given the ~9-month gestation lag: “AT&T iPhone” at the 2006/2007 boundary (the iPhone launched June 29, 2007) and “Sprint Android” at the 2009/2010 boundary (the HTC Hero launched October 11, 2009, so conceptions from that month onward yield births starting in mid-2010). The Sprint iPhone arrived October 14, 2011, after the conception window for every birth cohort in the sample, so a Sprint-iPhone effect cannot enter. Markers and confidence intervals are rendered in Sprint yellow. Bars are 95% confidence intervals from a county-clustered bootstrap with 500 replications. Sources: NCHS (2025), SEER (2026), NTIA (2010).

Appendix C: ATUS-CPS geographic partition and time-use event study

This appendix documents an attempt to extend the natural-experimental design of the body to ATUS time-use outcomes. The ATUS-CPS file identifies a respondent’s state for every survey year but applies the standard CPS public-use suppression rule to county of residence: county is revealed only if the cell clears the disclosure threshold, otherwise only the CBSA or the state-level metro/non-metro flag is observed.

C.1 ATUS-CPS cell-level coverage assignment

The most granular geography we can attach to an ATUS respondent is one of three kinds of “cell” that we construct from the joint distribution of `gestfips`, `gtco`, and `gtcbsa` and how we assign AT&T coverage:

1. *Identified counties* ($n = 404$): a single county whose FIPS code ATUS-CPS reveals. The cell’s AT&T coverage is the county’s own population-weighted share from the SBI 2010 snapshot.
2. *Balance of CBSA cells* ($n = 776$): for a CBSA where some but not all constituent counties are individually identified, the “balance” is the set of unidentified counties. By construction, any respondent whose county is suppressed but whose CBSA is identified must be in the balance, never in any of the CBSA’s identified counties. We assign the balance’s population-weighted mean coverage across its constituent counties.
3. *Balance of state cells*: the set of non-CBSA counties of a state. We assign the state’s population-weighted mean across its non-CBSA counties. In six states (Montana, Nebraska, North Dakota, Wyoming, Utah, and New Mexico) every non-CBSA county has AT&T coverage at or below 10%, so the balance-of-state cell is unambiguously a clean “control” assignment with no within-cell misclassification. In the remaining states the balance-of-state mean is interpretable only as a continuous coverage measure subject to within-cell measurement error in proportion to within-state non-CBSA coverage heterogeneity.

Figure C.1 maps the resulting cell partition, with each cell shaded by its population-weighted AT&T 3G coverage. County borders are dissolved within

cells: a balance-of-state cell appears as one contiguous polygon spanning all of a state’s non-CBSA counties, a balance-of-CBSA cell appears as one polygon spanning its constituent unidentified counties, and an identified county appears as its own polygon. Light-gray cells (the strict-six states’ non-CBSA balances) have essentially zero coverage; the AT&T blues shade the remaining range from 0–10% (light) to over 90% (dark). All non-CBSA territory is filled in, so the map can be read directly as the AT&T coverage value the ATUS analysis assigns to a respondent observed in each cell.

C.2 Time with friends: Poisson event study by age

We estimate a Poisson event study of ATUS minutes per day with friends (excluding episodes during work or class) on the cell-level AT&T 3G coverage variable described in C.1, separately by 5-year age band over 2005–2011. Formally,

$$\mathbb{E}[m_{ict} | \alpha_c, \gamma_t, A_c, X_i] = \exp\left(\alpha_c + \gamma_t + \sum_{y \neq 2007} \beta_y A_c \cdot \mathbf{1}[t = y] + X_i' \delta\right), \quad (2)$$

where m_{ict} is minutes per day with friends for respondent i residing in cell c in survey year $t \in \{2005, 2006, 2008, \dots, 2011\}$; $A_c \in [0, 1]$ is the cell’s December 2010 AT&T 3G coverage share, constant across years; α_c and γ_t are additive cell and survey-year fixed effects; and X_i is a vector of respondent-level controls. The event-time coefficients β_y trace the year-by-year effect of a counterfactual increase from zero to full AT&T coverage, with $y = 2007$ omitted as the pre-iPhone reference. We estimate by Poisson maximum likelihood with ATUS final weights and standard errors clustered at the cell level.

The three nested specifications differ in the composition of X_i . Specification 1 includes race/ethnicity (four categories), sex, and day-of-week fixed effects. Specification 2 adds single-year-of-age fixed effects. Specification 3 further adds married-or-partnered status, educational attainment (four categories), employment, household children under 5, household children 5–17, and family-income category fixed effects (16 levels plus a missing dummy). The sample ($N = 41,103$) is restricted to respondents non-missing on every Specification 3 control so the three specifications are strictly nested.

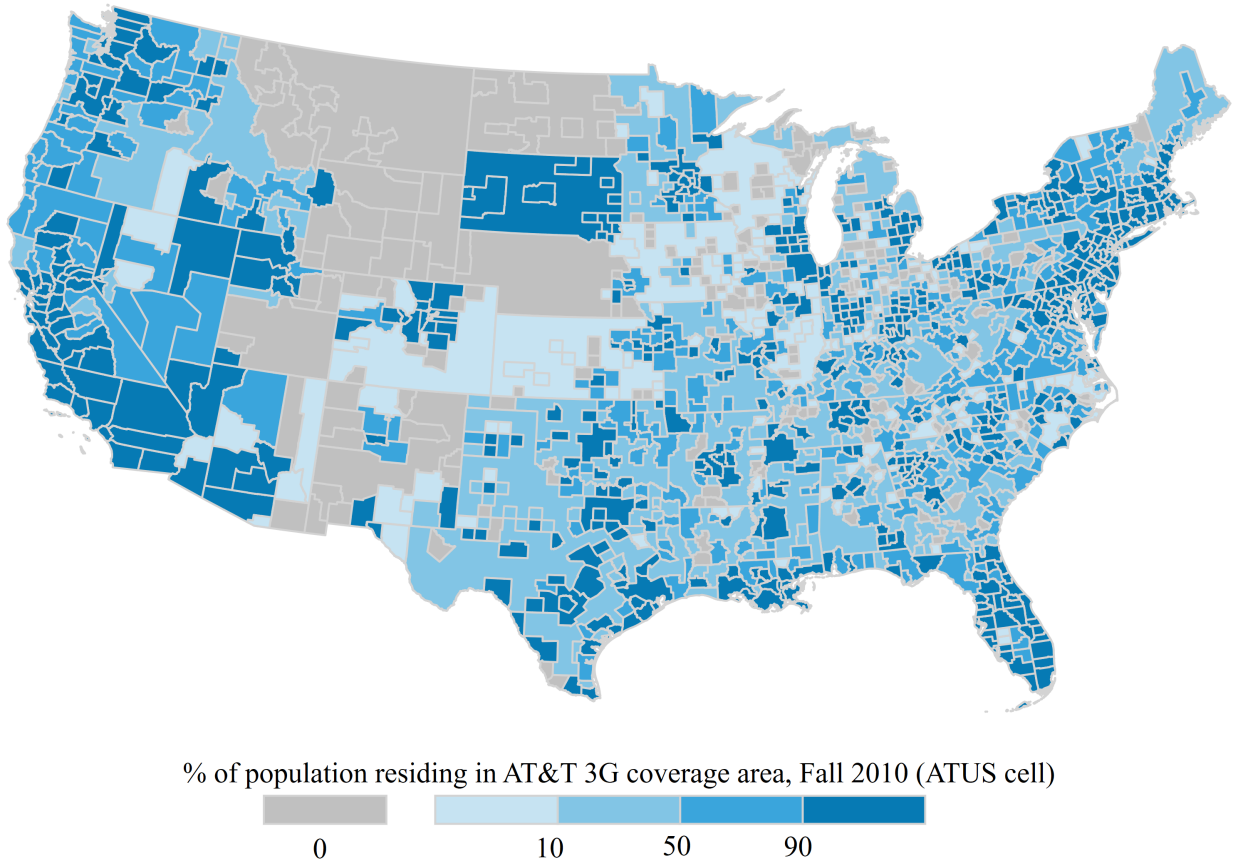
Figure C.2 plots the event-time coefficients. The Specification 2 post-2007 ATTs are negative at every age at which our headline fertility effects appear: –20.5%

at ages 15–19 ($p = 0.23$), -5.7% at 20–24 ($p = 0.85$), -12.4% at 25–29 ($p = 0.67$), and -27.7% at 30–34 ($p = 0.28$), where each percent change corresponds to the effect of a counterfactual increase from zero to full cell-level AT&T coverage. The point estimates are directionally consistent with the in-person-displacement channel, but every estimate is imprecise and statistically indistinguishable from zero at conventional levels; the 95% confidence intervals span both substantial declines and substantial increases at every age.

The imprecision reflects substantial measurement error on both sides of the regression. On the outcome side, the ATUS measure is a single 24-hour time diary per respondent. Time with friends on a single day is sharply zero-inflated and heavy-tailed—most respondents see no friends at all, and a small share spend several hours—so the individual outcome carries large within-cell variance even when averaged across thousands of respondents per cell-year. On the treatment side, the ATUS-CPS suppresses county for most respondents, so for the 60% of the sample outside the identified-county branch we assign coverage as a population-weighted mean across the counties in a balance-of-CBSA or balance-of-state cell. Where the constituent counties differ in their AT&T coverage, this cell-mean assignment introduces classical measurement error into the regressor, attenuating the point estimate toward zero. The combination—noisy outcome, attenuated coefficient—produces confidence intervals that are wide and centered closer to zero than the true effect likely lies.

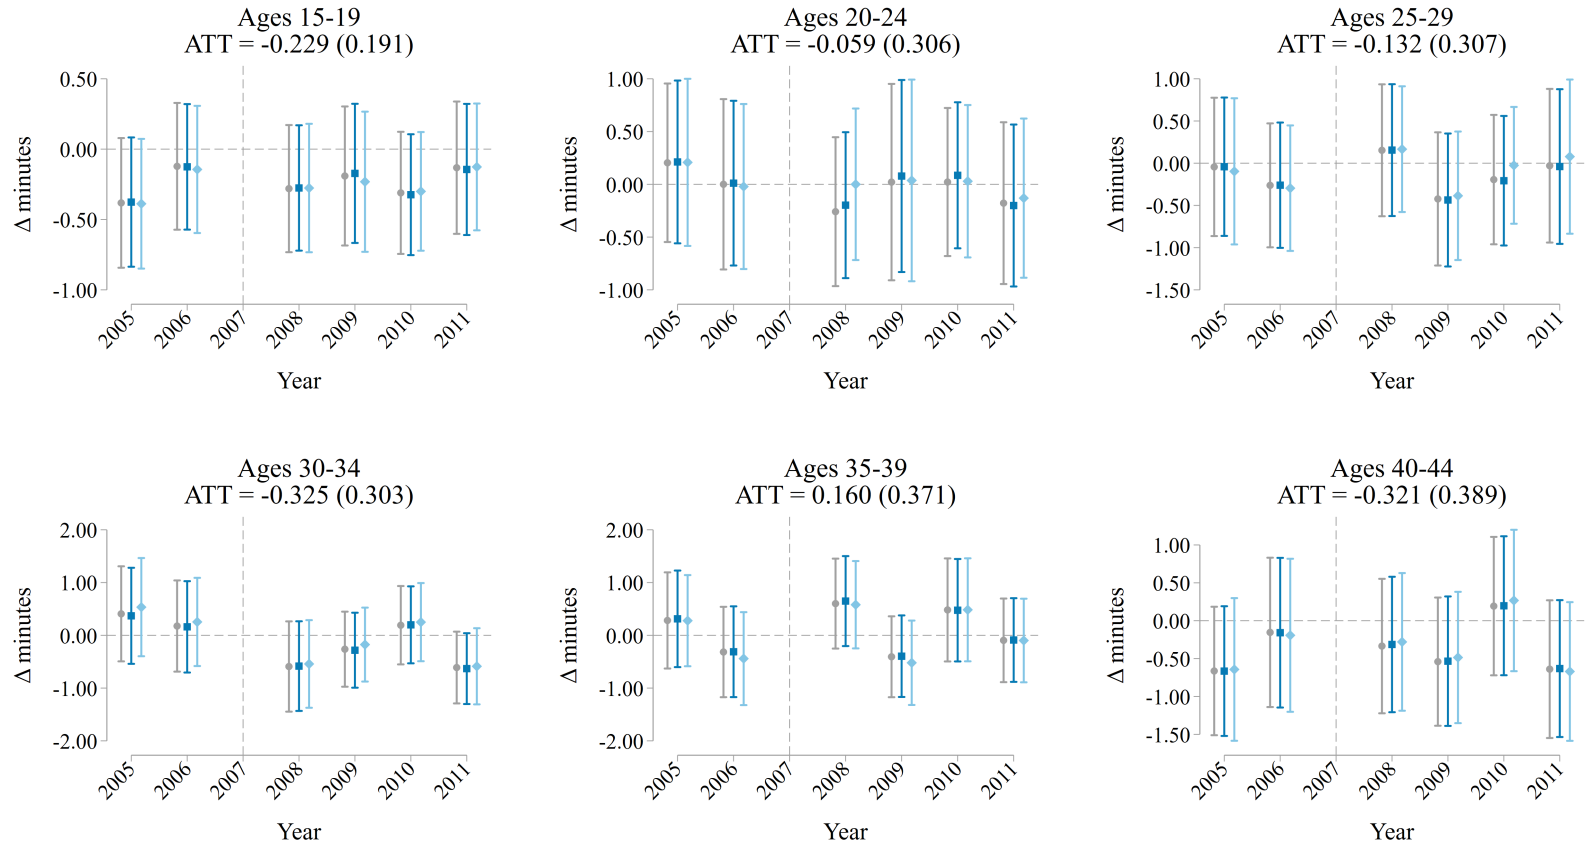
We therefore treat [Figure C.2](#) as suggestive rather than as a standalone causal estimate. The signs and rough magnitudes of the point estimates are consistent with the in-person-displacement channel that motivates the descriptive trends in [Figure 9](#), but the cell-level partition of the ATUS-CPS does not deliver the precision our main county-level natality design does.

Figure C.1: AT&T 3G coverage by ATUS-CPS assignment cell



Notes: Each polygon is one ATUS-CPS assignment cell: an identified county, the balance of a CBSA (its unidentified counties combined into a single polygon), or the balance of a state (its non-CBSA counties combined). Shading is the cell's population-weighted (women 15–44 in 2010) AT&T 3G coverage share (National Telecommunications and Information Administration (NTIA), 2010). Light gray (the lowest bin) shades cells with essentially zero coverage; the AT&T-blue gradient shades the remaining bins as in Figure 2. Alaska and Hawaii are omitted. The cell partition is described in the text above; the per-respondent assignment lives on `Output/Tables/ATUS.NonHHTime.Individual.dta` as `att_cov_cts` (continuous) and `att_cov_strict` (clean assignment only). Source: SBI (2010), Census 2013 CBSA delineation file, ATUS multi-year ATUS-CPS file (2025).

Figure C.2: Estimated effect of continuous AT&T 3G coverage on minutes spent in presence of friends, by age, from a Poisson event study under three nested control specifications (all ATUS cells)



Controls: ● Race + sex + day-of-week ■ + age FE ◆ + marital + educ + emp + kids + income

Notes: Each panel reports the event-time coefficients $\hat{\beta}_y$ from the Poisson (PPML) event study in equation (2), estimated on minutes per day with friends (excluding episodes during work or class) for the indicated 5-year age band. The sample is 41,103 ATUS respondents aged 15–44 in 2005–2011; 2003–2004 are dropped because ATUS-CPS suppresses CBSA and county identifiers in those years. The three nested control specifications match the legend; the full X_i ; enumeration, fixed-effect structure, and estimator are described in Appendix C.2. Coefficients have a Poisson log-rate interpretation; vertical bars are 95% confidence intervals from cell-clustered standard errors. The subtitle under each age label reports the post-2007 ATT (mean of the 2008–2011 coefficients) and its delta-method SE for Specification 2. Sources: ATUS (BLS, 2025); NTIA (2010).