

THE COLLAPSE OF TEEN FERTILITY IN THE DIGITAL ERA

NATHAN HUDSON* HERNAN MOSCOSO BOEDO*[†]

**Department of Economics, Carl H. Lindner College of Business, University of Cincinnati.*

April 25, 2026

ABSTRACT

Teen fertility collapsed globally starting around 2007. This affected countries across the income and policy spectrum. This paper argues that smartphones changed how teens spend time with each other, and that this change in turn drove the collapse in teen fertility. Once enough teens are on the phone, being on the phone is where the peer network is; in-person time falls sharply, and with it the unstructured contact in which most unintended teen conceptions occur. A coordination model formalizes this tipping: as the smartphone price falls, the in-person equilibrium ceases to exist and the economy moves to a phone-mediated one. Within the United States, terrain-ruggedness variation in broadband and 4G coverage identifies a causal effect on teen fertility, and time-use diaries show in-person socializing among teens roughly halving while digital leisure roughly tripled. A parallel design for England and Wales recovers the same acceleration and the same effect of mobile coverage on teen conceptions, ruling out country-specific contraceptive-access and welfare-reform stories. The model predicts that the shift towards the phone-mediated equilibrium affects multiple aspects of teen behavior. The same instrument that produces a collapse in teen fertility produces a surge in teen suicides.

Keywords: fertility, smartphones, teen pregnancy, broadband, coordination failure, instrumental variables, time use, international replication.

JEL Codes: J13, J11, J12, O33, L96, D85.

[†]Corresponding author: hernan.moscoso@uc.edu.

I. INTRODUCTION

Teen fertility collapsed almost everywhere in the world starting around 2007. This is a striking fact. Countries with very different healthcare systems, welfare regimes, abortion laws, religious traditions, recessions, and demographic trends all saw similar breaks in the same window. Whatever caused it was something global—something that arrived in roughly the same form in all of these places at roughly the same time.

This paper proposes a mechanism that runs through how teens spend time with each other. The mechanism is a coordination problem in peer-time allocation: each teen prefers to be wherever everyone else is, and once smartphone usage passes a threshold, the in-person equilibrium becomes self-undermining and is replaced by a phone-mediated one. The result is a discrete change in behavior—a tipping—that follows from a smooth underlying change in price. As the equilibrium shifts from in-person to digitally connected, a range of teen outcomes are affected. Sexual activity, fertility, suicide, violent crime, and depression all see changes in their dynamics in the same window, with signs that depend on whether each outcome rises or falls with peer presence. The paper’s empirical work is anchored on fertility, where the identification is sharpest, and uses suicide as a sign-flipping test of the equilibrium-shift mechanism.

The rest of the paper develops this argument and the evidence behind it. Figure 1 sets out the global pattern that motivates the analysis. Panel A plots the three-year centered log-linear growth rate of the teen birth rate (ages 15–19) across 128 countries. In the high-income subsample, the median growth rate moves sharply downward starting in 2008 and stabilizes several percentage points below its pre-2007 trajectory. Re-centering the time axis at each country’s own smartphone-shock onset—defined as $\tau_c = \max(\text{iPhone}_c, \text{mobile80}_c)$, the later of iPhone country-rollout and crossing 80 mobile cellular subscriptions per 100 people—the same break appears in the full 128-country sample.² Panel B turns from the headline fact to the age signature, plotting age-specific

²We treat the iPhone as a salient signal for the wider digital era rather than as the operative cause. The mechanism documented in this paper does not depend on Apple’s specific device: the timing of the smartphone shock in each country is the date at which the broader digital-infrastructure bundle—always-on connectivity, smartphone-grade handsets, and a mass-market app ecosystem—reached enough of the teen

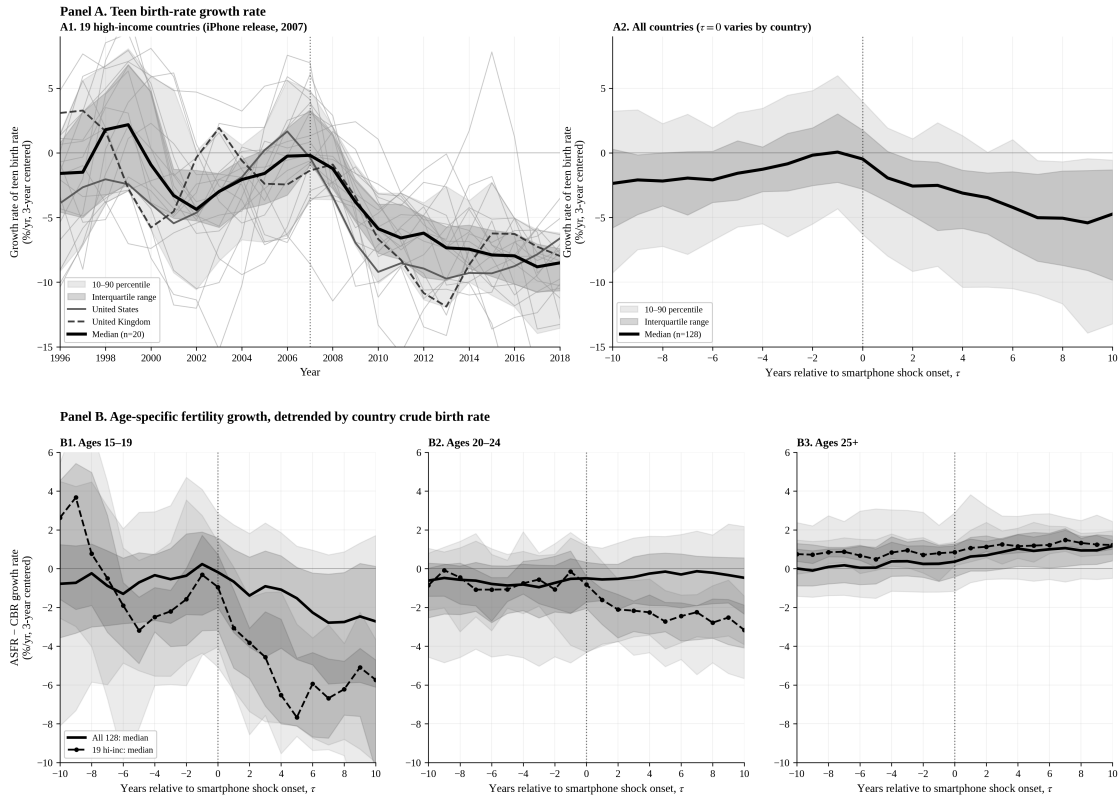


FIGURE 1. Teen fertility in the digital era. **Panel A** shows the three-year centered log-linear growth rate of the teen birth rate. Sub-panel A1: 19 high-income countries, 1996–2018, in calendar time; the dotted vertical line marks 2007, the iPhone release year. Sub-panel A2: all 128 countries, with time re-centered at each country's smartphone-shock onset $\tau_c = \max(\text{iPhone}_c, \text{mobile80}_c)$. **Panel B** shows age-specific fertility rate growth detrended by the country's own crude birth rate growth, recentered at τ_c , for three age aggregates: B1 ages 15–19, B2 ages 20–24, and B3 ages 25+ (population-weighted aggregate of 25–29, 30–34, 35–39, 40–44, 45–49). Heavy black line: full-sample median; dashed line with markers: 19 high-income countries' median; shaded bands: interquartile range (dark) and 10–90th percentile range (light), full sample only in Panel B. Sources: UN World Population Prospects 2024 (age-specific fertility rates, median variant); World Bank WDI (crude birth rate SP.DYN.CBRT.IN; mobile cellular subscriptions IT.CEL.SETS.P2); Apple press releases (iPhone country rollout dates).

fertility growth detrended by each country's own crude birth rate growth—a within-country detrending that nets out recessions, family policies, contraception access changes, and demographic-transition dynamics that move all ages together. The detrended deviation falls sharply below zero post-onset for ages 15–19, modestly for ages 20–24, and is essentially flat at zero for the 25+ aggregate. Whatever the smartphone shock did to fertility, it did to

peer network to tip the coordination equilibrium. We use iPhone country rollout dates because they are well-documented and tightly clustered with the wider smartphone-era bundle in each country.

teens; the 25+ population, which accounts for roughly 80 percent of women of reproductive age, exhibits no detrended response in the typical country.

This pattern is inconsistent with country-specific policy explanations: the relevant policies do not exist in Norway, Sweden, Switzerland, Luxembourg, Japan, or Australia, and Australia did not experience the 2008–2009 recession at all. The pattern is also inconsistent with the 2008 financial crisis as the proximate cause: the depth and timing of the recession varied sharply across the 19 high-income countries, but the teen birth rate break did not. The near-universal timing of the break across countries with fundamentally different healthcare, welfare, economic, and cultural environments points to a common global technology shock.

We use the phone as the key to the digital world: it is the device through which teen peer interaction migrates from in-person contact to digitally mediated contact, and its falling real price is the proximate trigger that tips the coordination equilibrium.³ The mechanism is strictly one-way: the hedonic quality-adjusted price of the phone has been falling steadily for years—by roughly 84 percent in real terms between 2007 and 2019 (Aizcorbe, Byrne, and Sichel, 2019), with substantial pre-2007 declines as well. As the price falls, more teens adopt the phone, and as more teens adopt, the value of being on the digital network grows through a standard network externality. At some point—identified in the data as 2006–2007, when U.S. teen cell-phone ownership crossed 50 percent between the November 2004 and fall 2006 Pew waves and the iPhone launch raised the salience of the smartphone category—the network became large enough that the externality triggered a cascade, and the coordination equilibrium jumped from the in-person to the digitally mediated side. The tipping date is not a feature of the price trajectory, which was smoothly

³A natural alternative hypothesis is that mobile-phone radiofrequency radiation directly impairs fertility rather than operating through behavior. Laboratory studies document reductions in sperm motility, viability, and count under phone exposure (Adams et al., 2014; Yu et al., 2021; Rahban et al., 2023), but at the population level no study has identified an effect on aggregate conception or birth rates, and the international human-health committees (ICNIRP, SCENIHR) draw no firm conclusions at typical exposure levels (Kenny et al., 2024). Two features of the data argue against a direct biological channel here. First, the age gradient runs the wrong way: the birth rate among women 15–19 fell by 71 percent while the rate among women 35–39 rose by 9 percent, even though phone use rose across all adult ages. Second, the decline is dominated by a collapse in conceptions rather than live births (Section 6.C), which is the opposite of what a fertility-impairment channel would predict.

declining; it is a feature of the endogenous adoption response, which reached critical mass at that moment. Peer-time allocation exhibits positive network externalities in both in-person and digitally mediated interaction, which generates multiple equilibria; the falling price pulls in enough of the teen peer network that eventually the digital-side externality dominates, with a mechanical reduction in conceptions as a consequence.

The U.S. case allows us to see the mechanism. Four pieces of evidence from the U.S. identify the channel. First, the age gradient operates with near-surgical precision: between 2007 and 2024, the birth rate among women 15–19 fell by 71 percent, 20–24 by 43 percent, 25–29 by 23 percent, 30–34 by 1 percent, and 35–39 *rose* by 9 percent. The monotone gradient with a zero in the early thirties is difficult to rationalize with any candidate broad economic force but fits naturally with a mechanism specific to the unstructured peer-contact environment of adolescents and young adults. Second, the decline is neither tempo nor a shift in pregnancy resolution: cohort data rule out delay (women born in 1990 have completed only 1.62 children by age 33 vs. 2.16 for the 1975 cohort), and Guttmacher data rule out rising terminations (among women 15–17 over 2007–2020, pregnancy fell 70 percent and the abortion ratio held stable at 29–30 percent). What fell was conceptions. Third, two complementary county-level specifications view the same phenomenon through different lenses. A cross-sectional IV using terrain ruggedness (Akerman, Gaarder, and Mogstad, 2015) as instrument for digital-infrastructure adoption, run in parallel for fixed broadband and 4G LTE, recovers significant causal effects with first-stage F -statistics well above conventional thresholds. A within-county distributed-lag panel of 3,135 counties yields gradually-building four-year cumulative effects on both measures, with Engle-Granger cointegration confirming a stable long-run broadband relationship at a 1.7-year half-life. Fourth, ATUS diary data measure the behavioral mechanism directly: U.S. teens reduced in-person socializing from 68 minutes per day in 2003 to 38 in 2019 while computer-leisure rose from 22 to 96 minutes; structured activities like sports were preserved while unstructured peer time was displaced— exactly the pattern the model predicts.

The English and Welsh case confirms the mechanism. The paper replicates the U.S.

identification strategy for 294 local authorities in England and Wales. ONS publishes under-18 conception rates (a direct measure of the theoretical conception hazard) by local authority annually 1998–2021, and Ofcom Connected Nations data provide local-authority 4G coverage. Because the NHS provides universal contraception, the welfare-reform and Medicaid-expansion channels are absent in this setting. Under-18 conception rates accelerated from -1.3 percent per year pre-2007 to -8.9 percent post-2007, a $6.8\times$ acceleration that nearly matches the U.S. $5.6\times$ acceleration.⁴ Recovering the same coefficients in two countries with completely different policy environments argues strongly that the channel is not a U.S.-specific artifact and rules out pharmaceutical-access and welfare-reform explanations.

The paper’s central framing is therefore that smartphones *accelerated* a decline already underway for reasons documented in prior literature, rather than *caused* the decline ab initio. The 1991–2005 contraceptive-access story remains correct for that period. The paper’s contribution is to identify a specific additional force—a peer-network coordination channel generating a hysteretic regime change—that explains why the decline accelerated to roughly double its prior pace after 2007.

The mechanism’s distinctive prediction. The coordination model of Section 3 makes a sharper empirical prediction than a smooth-transition story would. In a smooth-transition account the falling smartphone price generates a continuous behavioral response, with no reason for the second derivative of any outcome to change at any particular date. In the coordination-tipping account two equilibria coexist until the price falls below a critical threshold, at which point the in-person equilibrium ceases to exist and the peer network jumps to the digital equilibrium. Outcomes that depend on in-person peer time should therefore exhibit synchronized *kinks*—changes in slope, not changes in level—at the same tipping date, with the sign of each kink determined by whether the outcome rises

⁴The $5.6\times$ and $6.8\times$ ratios compare the post-2007 annualized rate of decline to the pre-2007 annualized rate of decline within each country. This is a different quantity than the U.S. pre-/post- ratio computed on teen pregnancy rates (ages 15–17), which is roughly $2\times$ (4.9% to 9.8% annualized decline; see Section 6.1). The two ratios use different outcome series (UK conceptions vs. U.S. pregnancies; ages 15–19 vs. 15–17) over different windows; the headline finding is that both U.S. and UK declines steepened sharply post-2007, by similar proportional amounts in each country.

or falls with peer presence. Section 8 tests this prediction. Five U.S. teen behavioral outcomes—the teen birth rate, sexual activity, violent-crime arrests, depression, and the suicide rate—all kink at 2007 in the directions the model predicts, with peer-time-dependent outcomes accelerating downward and isolation-correlated outcomes reversing upward. The international teen suicide rate displays the same U-shape across the 19 high-income countries used in the cross-country fertility analysis. Section 8 also extends the within-country causal identification of Section 6 to the suicide margin: applying the same terrain-ruggedness instrument to teen suicide rates on a 1,899-county sample (recovered through an all-age aggregate decomposition that circumvents CDC’s small-cell suppression rule), we obtain a positive effect of 4G coverage on the post-period teen suicide rate—a 10-percentage-point increase in coverage is associated with roughly 2.7 to 4.0 additional teen suicide deaths per 100,000 depending on specification, against a sample-mean post-period rate of 13.9 per 100,000. The estimate is opposite-signed to the fertility result and of similar proportional order of magnitude, exactly as the model predicts when the same shock acts on outcomes that rise or fall with peer presence. The broader behavioral patterns we discuss in Section 8 both confirm the model’s distinctive empirical signature and complement causal evidence on the mental-health margin developed independently in the psychology literature (Twenge, Martin, and Campbell, 2018; Twenge et al., 2019; Haidt, 2024) and in quasi-experimental economics (Braghieri, Levy, and Makarin, 2022; Donati et al., 2025).

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 presents a coordination-equilibrium model of teen peer time and its fertility implications. Section 4 describes the data. Section 5 develops the cross-country descriptive evidence in detail. Section 6 presents the within-country causal identification for the United States. Section 7 replicates the analysis for England and Wales. Section 8 tests the model’s distinctive prediction—a synchronized kink across peer-time-dependent outcomes at the tipping date—against U.S. and international data on teen behaviors beyond fertility. Section 9 concludes.

II. RELATED LITERATURE

This paper connects four bodies of work: the macroeconomics of recent fertility decline, the technology-and-behavior literature on smartphones and social media, the demographic literature on teen pregnancy, and the game-theoretic literature on network externalities and coordination failure.

Kearney, Levine, and Pardue (2022) document the puzzle in its starkest form: the fall in the U.S. total fertility rate since the Great Recession cannot be explained by unemployment, housing costs, child-care costs, or student debt, and appears to reflect a shift in the desire to have children rather than in the capacity to do so. Doepke et al. (2023) emphasize that the post-2007 decline occurs simultaneously across countries with very different labor-market and policy environments, ruling out most country-specific explanations. We propose the smartphone as a specific common technology shock operating through the peer-coordination environment in which young adults have unstructured in-person contact.

Allcott et al. (2020) use experimental Facebook deactivation to estimate welfare effects; Braghieri, Levy, and Makarin (2022) exploit the staggered Facebook rollout to identify mental-health effects; and Twenge, Martin, and Campbell (2018); Haidt (2024) document time-series co-movements between smartphone diffusion and adolescent well-being. Our contribution is to apply the cross-sectional smartphone-adoption identification strategy to fertility rather than mental health.

The teen-pregnancy literature has long been dominated by contraceptive-access and welfare-reform explanations. Kearney and Levine (2015) attribute the 1991–2008 decline roughly equally to contraceptive improvements and delayed sexual initiation, with these forces approaching saturation by 2005. We take those conclusions as settled for the pre-2005 period; the post-2007 re-acceleration is a separate phenomenon driven by the smartphone diffusion documented here.

The theoretical model builds on the game-theoretic tradition of coordination failure (Schelling, 1978; Milgrom and Roberts, 1990) and on the durable-goods-and-time tradition that rationalizes how falling hedonic prices reorganize household time use (Greenwood,

Seshadri, and Yorukoglu, 2005; Manuelli and Seshadri, 2014). The smartphone case is distinguished by a coordination externality in peer-time allocation: a smartphone is valuable only to the extent one’s peers are on the platform, and in-person socializing is valuable only to the extent peers are physically available. The resulting strategic complementarity generates the tipping dynamic absent from appliance or tractor diffusion paths. We therefore adopt a minimal static coordination model rather than a quantitative general equilibrium structure: the behavior under study is a teen’s intratemporal allocation of a specific afternoon’s peer time, with no saving, borrowing, market-clearing prices, or intertemporal optimization.

III. A COORDINATION MODEL WITH A FALLING PHONE PRICE

This section presents a minimal coordination model of peer-time allocation in which the only exogenous shock is the falling hedonic price of the smartphone. The model delivers three results that map directly onto the empirical facts: multiple equilibria separated by a tipping point, a regime change triggered by the price decline, and hysteresis once the new equilibrium is reached.

III.A Environment

Consider a population of young individuals indexed by $i \in [0, 1]$. During unstructured peer time each individual chooses one interaction medium: in-person co-location ($a_i = 0$) or phone-mediated contact ($a_i = 1$). Let $m \equiv \int a_i di$ denote the aggregate share of the peer network on the phone.

The payoff to in-person contact is

$$u_{\text{in}}(m, \theta_i) = \alpha(1 - m) + \theta_i, \tag{1}$$

where $\alpha > 0$ is the value of being in-person when peers are also there, and $\theta_i \sim F(\theta)$ is an idiosyncratic preference shock capturing heterogeneity in the taste for in-person socializing (extroversion, local social capital, access to transportation). The payoff to phone-mediated

contact is

$$u_{\text{phone}}(m; p) = \beta + \delta m - p, \quad (2)$$

where $\beta \geq 0$ is baseline device utility, $\delta > 0$ is the phone network externality, and $p \geq 0$ is the quality-adjusted phone price.

The key symmetry is that both activities exhibit positive own-network externalities: in-person is more valuable when others are in person, and phone is more valuable when others are on phones. This is the standard supermodular-game structure of Katz and Shapiro (1985) and Milgrom and Roberts (1990), but applied here to the medium of peer communication rather than to technology adoption per se.

III.B Multiple Equilibria and the Price Shock

Individual i picks the phone ($a_i = 1$) if and only if $u_{\text{phone}}(m; p) \geq u_{\text{in}}(m, \theta_i)$, which simplifies to

$$\theta_i \leq \bar{\theta}(m; p) \equiv \beta - p - \alpha + (\delta + \alpha)m.$$

Aggregating over the population, a symmetric Bayesian-Nash equilibrium is a fixed point $m^* = F(\bar{\theta}(m^*; p))$.

Proposition 1 (Multiple Equilibria). *Suppose the preference distribution F has a continuous density f satisfying*

$$(\delta + \alpha) \sup_{\theta \in \mathbb{R}} f(\theta) > 1. \quad (3)$$

Then there exist prices $\underline{p} < \bar{p}$ such that:

1. *For every $p \in [\underline{p}, \bar{p}]$ the equilibrium correspondence $m^* = F(\bar{\theta}(m^*; p))$ admits exactly three fixed points: a stable low-adoption equilibrium $m_L^*(p)$, an unstable interior fixed point $m_M^*(p)$, and a stable high-adoption equilibrium $m_H^*(p)$, with $m_L^*(p) < m_M^*(p) < m_H^*(p)$.*
2. *For every $p \notin [\underline{p}, \bar{p}]$, the fixed point is unique.*

Stability is in the sense of best-response dynamics: $\partial G(m; p) / \partial m|_{m=m^} < 1$ at m_L^* and m_H^* ,*

and > 1 at m_M^* , where $G(m; p) \equiv F(\bar{\theta}(m; p))$.⁵

Condition (3) is structural: it depends only on the primitives F , δ , and α , not on the price p . Economically, it says network externalities ($\delta + \alpha$) are large relative to the spread of preferences—peer coordination matters a lot in how young people socialize, which has been well-documented since Schelling (1978). The price range $[p, \bar{p}]$ is not an additional assumption but a consequence of the structural condition: as p falls smoothly through the window, the equilibrium count goes $1 \rightarrow 3 \rightarrow 1$, generating the discrete behavioral jump that motivates the empirical analysis. Figure 2 illustrates the three regimes.

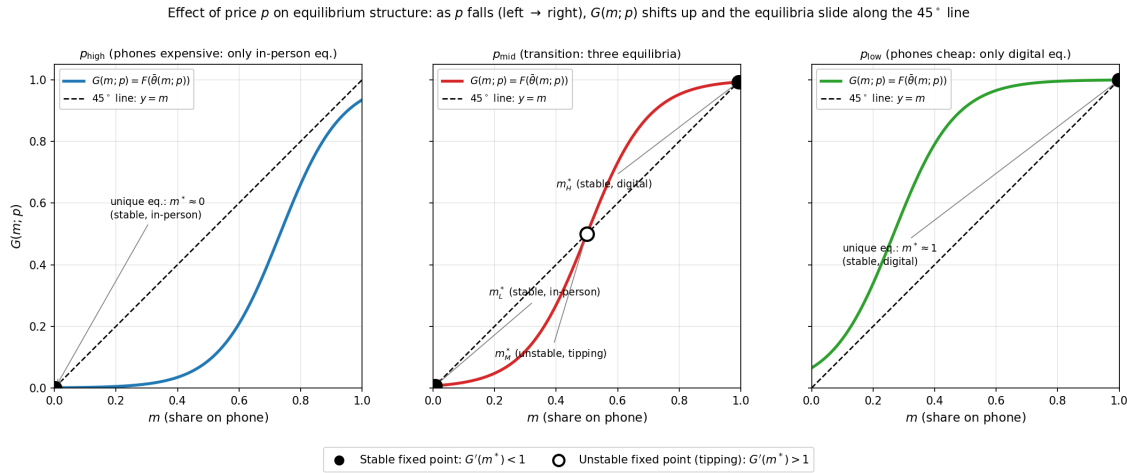


FIGURE 2. Effect of price p on the equilibrium structure. Each panel plots $G(m; p) \equiv F(\bar{\theta}(m; p))$ (colored S-curve) against the 45° line $y = m$ (dashed); fixed points are intersections. Filled dots are stable; the open dot is unstable. Left: high p , unique in-person equilibrium. Middle: $p \in [p, \bar{p}]$, three coexisting equilibria. Right: low p , unique digital equilibrium. Parameters: $\alpha = 1$, $\beta = 1.5$, $\delta = 2$; F logistic with scale 0.30, satisfying $(\delta + \alpha) \sup f = 2.5 > 1$.

The price as the single exogenous mover. In the model, the only exogenous object that moves is the phone price p . Adoption m is endogenous: it responds to p through the threshold decision of equation (1). Figure 3, Panel A, plots a hedonic quality-adjusted smartphone price index constructed following the methodology of Aizcorbe, Byrne, and

⁵*Proof sketch.* Define $H(m; p) = G(m; p) - m$. Differentiating, $H'(m; p) = (\delta + \alpha)f(\bar{\theta}(m; p)) - 1$. Under (3), H' takes both positive and negative values as m ranges over $[0, 1]$, so H is non-monotone. For p in a non-empty interval $[p, \bar{p}]$ the non-monotonicity is positioned so that H changes sign three times—equivalently, G crosses the 45° line three times—delivering three fixed points. Outside this interval, $G(\cdot; p)$ has shifted far enough that H has only one sign change. Stability of m_L^* and m_H^* and instability of m_M^* follow from $H'(m^*; p) \leq 0$ at the three roots, equivalent to $G'(m^*; p) \leq 1$.

Sichel (2019). In the Aizcorbe-Byrne-Sichel series, the index declines at an average annual rate of 17 percent from 2010 to 2018. Cumulatively, from 2007 to 2019, the quality-adjusted price of a smartphone fell roughly 84 percent in real terms, and the series had already been falling for several years before 2007. This decline reflects falling hardware costs, rising processor and display performance per dollar, and growing carrier subsidies. The price trajectory is smooth. What is not smooth is the endogenous adoption response: as p falls past successive thresholds of the preference distribution $F(\theta)$, adoption m accelerates, and once m grows large enough the network externality δm on the phone side exceeds the in-person externality $\alpha(1 - m)$ for the marginal agent, triggering a cascade and an equilibrium switch. The tipping date is a feature of where the endogenous response function crosses critical mass, not a feature of any kink or discontinuity in p itself.

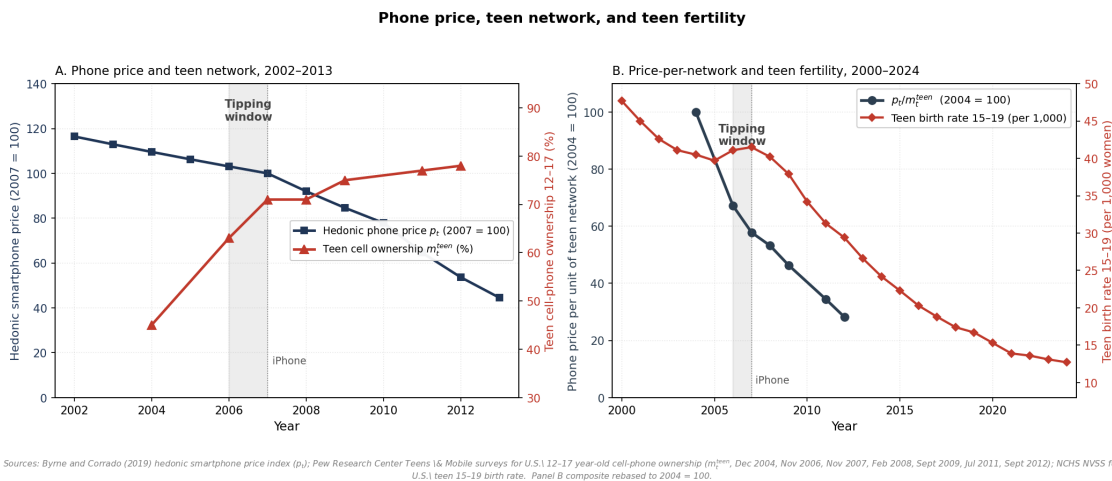


FIGURE 3. Phone price, teen network, and teen fertility. Panel A: Aizcorbe-Byrne-Sichel hedonic smartphone price index p_t (2007 = 100) and Pew Research Center cell-phone ownership among U.S. 12–17 year-olds m_t^{teen} , 2002–2013. The two series cross during the shaded tipping window (2006–2007). Panel B: the composite price-per-network p_t/m_t^{teen} (rebased to 2004 = 100) and the U.S. 15–19 teen birth rate, 2000–2024. The composite falls earliest and fastest through the tipping window; the teen birth rate follows through the conception channel with a lag consistent with the nine-month biological gap from conception to birth plus the time required for the new equilibrium to propagate through sexually experienced cohorts. The composite series ends in 2012 because Pew’s Teens & Mobile survey transitioned from cell-phone ownership to smartphone access beginning in 2013, a measurement discontinuity we do not bridge. Sources: Aizcorbe, Byrne, and Sichel (2019) for p_t ; Pew Research Center Teens & Mobile surveys (Nov 2004, Fall 2006, Nov 2007, Feb 2008, Sept 2009, Jul 2011, Sept 2012) for m_t^{teen} , with no interpolation between survey waves; NCHS NVSS for the teen birth rate.

In the model, the effect of falling p is to shift the threshold $\bar{\theta}(m; p)$ upward uniformly in m . For high p (the pre-smartphone regime), the equilibrium correspondence has only one fixed point at low m ; the phone equilibrium exists mathematically but is never reached because no individual will pay the high price to join an empty network. As p falls into the multiple-equilibrium region $[\underline{p}, \bar{p}]$, inertia keeps the system at m_L^* . Once p falls below \underline{p} , the low- m equilibrium disappears and the system must jump to the phone equilibrium m_H^* .

The timing implication matches the empirical record. The price trajectory contains no sharp feature at 2006–2007; the index falls smoothly through the mid-2000s. The sharp feature is in the endogenous adoption response. By the November 2004 Pew wave, 45 percent of U.S. teens owned a cell phone; by fall 2006, 63 percent. Crossing the 50 percent threshold placed the majority of the teen peer network on the phone side of the coordination problem, triggering the cascade. The iPhone launch of June 2007 is a realization of the continuous fall in p , not a separate shock; it accelerated the cascade but did not cause it ab initio. Panel B of Figure 3 summarizes both margins in the composite p_t/m_t^{teen} , which falls from 100 in 2004 to 28 by 2012. The fertility response develops over a longer horizon, from the 2006–2007 conception break through the deepening of the decline during 2010–2015 as successive cohorts entered the 15–24 age window with more phone-mediated peer exposure. Earlier technologies (landlines, pagers) never pulled enough of the network onto the phone side because p remained too high; the cell phone is distinctive because its quality-adjusted price fell far enough, fast enough, to push teen ownership past the critical mass by 2005–2006.

Proposition 2 (Hysteresis). *Let the economy be at $m_H^*(p_1)$ for some low $p_1 < \underline{p}$. If the price subsequently rises to $p_2 \in [\underline{p}, \bar{p}]$ where both equilibria exist, the economy remains at $m_H^*(p_2)$ rather than returning to $m_L^*(p_2)$.*

The hysteresis result has direct policy content: removing individual smartphones (raising p for one agent, or for a restricted subset) leaves aggregate m essentially unchanged, so the restricted agent faces $m \approx m_H^*$ and cannot unilaterally reallocate time to in-person uses. A coordinated shift of the peer network is required.

III.C Fertility Consequences

Assume conceptions arise only during in-person peer time. The per-period conception hazard for a young woman is

$$\lambda(m) = \lambda_0 \cdot (1 - m), \quad (4)$$

where λ_0 is the hazard when the economy is fully in person. As smartphones diffuse and m rises from m_L^* to m_H^* , the conception hazard falls proportionally. The 44-percent reduction in teen in-person socializing documented in the ATUS data of Section 6.F (Figure 8) provides direct empirical support for the reduced-form implication of equation (4).

The age gradient is mechanical. A cohort that enters the phone equilibrium at an early age spends a larger fraction of its reproductive years at the elevated m , and therefore loses more of its lifetime conception potential. A cohort that was already completing its family when the tipping occurred is affected only during its residual reproductive years. Because age-specific conception intensity $\lambda_0(y)$ peaks in the late teens and early twenties, the cohorts that lose the most are those who were adolescents when p fell through the tipping region—exactly the 15–24 age band that Figure 4 shows declining fastest.

IV. DATA

The empirical analysis draws on ten public data sources.

National age-specific birth rates are from NCHS natality files, annually 1990–2023. State-level teen birth rates (ages 15–19) are from the NCHS Data Visualization Gallery covering 1990–2019. County-level teen birth rates (ages 15–19, 2003–2020) are Khan-Rossen hierarchical-Bayesian smoothed estimates from the NCHS county-level teen births file (Khan et al., 2024), providing stable rates for 3,136 counties even in small populations.

Pregnancy and abortion rates by age group are from the Guttmacher Institute (Chiu, Maddow-Zimet, and Kost, 2024). The 2020 endpoint predates *Dobbs* so the data measure the pregnancy-abortion relationship in a constitutionally protected environment.

Broadband infrastructure data at the county level are from FCC Form 477. Two files cover 2008–2024: the historical county-tier file (December 2008 through December 2013) and the current county-tier file (June 2014 through December 2024). Each file reports, for each county and time point, a discrete tier code (0–5) indicating household residential broadband connection intensity per 1,000 households at successive speed thresholds. The speed-threshold definitions changed between the historical and current files; we stitch them together at the 10 Mbps/1 Mbps threshold (historical Tier 4, current Tier 2) to construct a consistent 2008–2024 measure of fast-broadband adoption at the county level. The stitched series has correlation 0.89 across the 2013–2014 break, validating continuity. Broadband is a proxy for smartphone-enabling infrastructure: carriers deployed fiber backhaul to cell towers along the same routes as cable broadband to households, so county-level variation in fixed broadband closely tracks county-level variation in mobile network readiness.

Mobile LTE (4G) coverage data at the county level are from FCC mobile deployment reports 2010–2019. For each county-year we observe the percent of premises with outdoor LTE coverage from all four national mobile network operators (Verizon, AT&T, T-Mobile, Sprint), which we take as a direct measure of mobile-network availability in the county. The LTE data offer a tighter proxy for the infrastructure that actually carries phone-mediated peer interaction. For the within-county panel analysis of Section 6.6, we extend the LTE series back to 2003 by imposing zero-fill for 2003–2009: this is a genuine zero because commercial LTE did not begin in the U.S. until Verizon’s December 2010 launch. The resulting 18-year series (2003–2020, with 2019 carried to 2020) provides annual county-level 4G coverage alongside the broadband tier in a consistent panel. We use LTE coverage as the treatment variable in the parallel cross-sectional IV specifications of Table 3, columns (4)–(6), and in the parallel within-county panel specifications of Tables 4 and 5.

Smartphone ownership by age group (national-level series) is from the Pew Research Center Mobile Fact Sheet (2011–2025), with teen access from Teens & Tech reports.

Behavioral indicators are from YRBS (CDC), Monitoring the Future (MTF), and FHWA Highway Statistics.

Cohort completed fertility is constructed from NCHS age-specific period birth rates, cumulating the observed age- x rate in year $t = c + x$ for cohort c .

County-level religious composition is from the 2010 U.S. Religion Census (Grammich et al., 2012), which enumerates adherents of 236 religious groups in every U.S. county. We use the four aggregated traditions—Evangelical Protestant, Catholic, Mainline Protestant, and Black Protestant—expressed as shares of county population. County-level political composition is from the MIT Election Data and Science Lab (MIT Election Data and Science Lab, 2024) county presidential returns file, which reports candidate-level vote totals for every U.S. county in the 2000, 2004, 2008, 2012, 2016, and 2020 presidential elections. We compute the county-level mean Republican share of the two-party vote across these six cycles as a measure of structural partisan composition. Both are used only as robustness controls in the county IV specifications of Table 3, columns (2), (3), (5), and (6).

For the cross-country analysis of Section 5, four additional sources are used. Country-year age-specific fertility rates by five-year age group (births per 1,000 women in each age group, ages 15–19 through 45–49) for all countries 1990–2023 are from the United Nations Department of Economic and Social Affairs, *World Population Prospects 2024*, medium variant (file WPP2024_Fertility_by_Age5). This source provides age-disaggregated fertility for the entire panel and underlies both the teen birth rate series in Panel A of Figure 1 and the age-decomposition series in Panel B. Country-year crude birth rates (births per 1,000 population) for the same panel are from the World Bank World Development Indicators series SP.DYN.CBRT.IN; we use the crude birth rate to detrend each country’s age-specific growth rates in Panel B of Figure 1, isolating the age signature of the smartphone shock from country-level fertility movements common to all ages. Country-year mobile cellular subscriptions per 100 people (2000–2023) are from the World Bank WDI series IT.CEL.SETS.P2, with underlying data from the ITU World Telecommunication/ICT Indicators Database. iPhone country rollout dates are hand-compiled from Apple’s official press releases: the original iPhone launch (United States, June 2007), the November 2007 launches in the United Kingdom, Germany, and France, the March 2008 launches in

Austria and Ireland, the July 11, 2008 iPhone 3G launch across 21 additional countries, the July 17, 2008 launch in France (iPhone 3G), and the August 22 and September 26, 2008 waves that completed Apple’s 2008 expansion to 70 countries. For countries that received the iPhone after 2008, launch years are from Wikipedia’s “iPhone 3G” and “iPhone 3GS” articles cross-referenced against Apple press releases and local carrier announcements.

For the suicide analysis of Section 8, county-level death counts and population denominators are from the CDC WONDER Multiple Cause of Death database, which provides death-record extracts by underlying cause (ICD-10), age group, and county for 1999–2020. WONDER suppresses any county-cell with fewer than ten deaths, so direct queries of teen 15–19 suicide deaths return reportable values for only about 360 of the 3,084 counties in the IV sample. Section 8.1 develops a decomposition that recovers teen suicide counts from queries on larger cells—all-ages all-external causes, all-ages non-suicide external causes, and not-15–19 suicide—each of which clears the ten-death threshold in nearly every county and together pin down teen suicide as a residual. This expands the recoverable sample from 363 counties to 1,899 counties spanning the rural–urban distribution.

V. CROSS-COUNTRY EVIDENCE

This section develops the cross-country descriptive evidence introduced in Figure 1 of the introduction. The logic is straightforward: if the post-2007 acceleration in teen birth rate decline is attributable to smartphone diffusion, then teen birth rate trajectories in other countries should display the same break pattern centered at each country’s own smartphone-shock onset rather than at the U.S. 2007 calendar year. Conversely, if the acceleration is instead attributable to some U.S.-specific policy or economic shock—the Affordable Care Act contraceptive mandate, the Colorado LARC initiative, Medicaid family-planning waivers, the 2008–2009 recession—no such pattern should appear in countries where those shocks did not occur. The section establishes that the teen birth rate break is global and tracks smartphone diffusion, not U.S.-specific policy.

V.A Treatment Variable and Sample

For each country c , the treatment year is

$$\tau_c = \max(\text{iPhone}_c, \text{mobile80}_c), \quad (5)$$

where iPhone_c is the year the iPhone first became commercially available in country c and mobile80_c is the year country c first crossed 80 mobile cellular subscriptions per 100 people. The max operator captures both preconditions for the smartphone shock: handset availability (Apple’s rollout) and feature-phone saturation, beyond which the marginal mobile adoption is a smartphone replacing a feature phone rather than a first phone.

For the 19 high-income countries, the iPhone constraint binds and τ_c clusters in 2007–2008. For the 128-country full sample with at least five annual observations pre- and post-treatment, the iPhone binds for the majority and the mobile-penetration threshold binds for the remainder (mostly low-income countries in sub-Saharan Africa, South Asia, and Latin America). Treatment years span 2007–2018.

V.B Growth Rate Construction

The outcome is the three-year centered log-linear growth rate of the teen birth rate, computed country-by-country as the OLS slope of $\log(BR_{c,t})$ on t over $\{t-1, t, t+1\}$, multiplied by 100. The event-time axis is $\tau_{c,t} = t - \tau_c$.

V.C Results

Panel A of Figure 1 presents the central result. Three features are worth calling out.

The break is sharp and well-timed. In the 19-country sub-panel, the median growth rate is -0.3 percent per year in 2007, falls to -1.5 in 2008, -3.6 in 2009, -5.7 in 2010, and stabilizes between -6 and -8 percent through 2017. The break sits one year after the 2007 U.S. iPhone launch and coincides with the July 2008 iPhone 3G global rollout. In the full 128-country sub-panel, the median is flat near -2 percent per year through $\tau = -2$, then drifts to approximately -5 percent by $\tau = +10$. Cross-country dispersion

also compresses post-treatment: the IQR narrows from 6 to 4 percentage points by $\tau = +5$.

The pattern is not confined to high-income countries. Table 1 decomposes the pre–post change by binding constraint. The 19 high-income countries show a change of -4.68 percentage points per year; the 98 iPhone-binding countries (high- plus middle-income) show -2.85 ; and the 33 mobile-threshold-binding (mostly low-income) countries show a smaller but still negative -1.23 . The monotone attenuation as the sample moves from high- to low-income is consistent with the coordination mechanism operating most strongly where the in-person peer equilibrium was firmly established.

TABLE 1. Pre-post median growth rate change by binding-constraint subsample

Subsample	n	Pre \bar{g} $\tau \in [-5, -2]$	Post \bar{g} $\tau \in [3, 6]$	Δ
19 high-income	19	-2.17	-6.85	-4.68
All countries (full sample)	131	-0.90	-3.34	-2.44
iPhone binds	98	-0.82	-3.67	-2.85
80% mobile binds	33	-1.22	-2.45	-1.23

Cross-country median three-year centered log-linear growth rate of the teen birth rate (percent per year). Pre window averages event times $\tau \in \{-5, -4, -3, -2\}$; post window $\tau \in \{3, 4, 5, 6\}$.

Robustness. Varying the mobile-penetration threshold across 50, 60, 70, 80, and 90 percent leaves the pre-post change at -4.68 to -5.08 for the 19-country subsample and -2.01 to -2.84 for the full sample. The 80 percent baseline matches the telecom-industry convention of feature-phone saturation (GSMA Intelligence, 2010), but the result does not hinge on it.

V.D Age Decomposition: Isolating the Smartphone Signature

The cross-country break in the teen birth rate documented above is striking, but it leaves open a question: is the post-2007 fertility decline a teen-specific phenomenon or a broader fertility movement that happens to show up most clearly at the bottom of the age distribution? A coordination-mechanism story in which smartphones reallocate teen peer time from in-person to phone-mediated interaction predicts an age-targeted response: the smartphone shock should disproportionately affect the age group whose

unintended-conception rate depends most heavily on unstructured peer time. A general-fertility story—a recession, a contraceptive-technology shock, a global cultural shift in desired family size—predicts a more uniform response across age groups, modulated by each age group’s typical contribution to total fertility.

Panel B of Figure 1 discriminates between these two hypotheses. For each country c , year t , and five-year age group a , we compute the three-year centered log-linear growth rate of that country’s age-specific fertility rate, $g_{c,a,t}^{\text{ASFR}}$, and subtract that country’s own three-year centered log-linear growth rate of the crude birth rate, $g_{c,t}^{\text{CBR}}$:

$$\text{rel}_{c,a,t} = g_{c,a,t}^{\text{ASFR}} - g_{c,t}^{\text{CBR}}. \quad (6)$$

This within-country detrending nets out anything that moves all ages of country c ’s fertility together—recessions, family policies, contraception-access changes, demographic-transition dynamics, cohort-size shifts, secular fertility decline—and leaves only the age-pattern signature of whatever shock is hitting different cohorts differently. We aggregate ages 25 and older into a single population-weighted measure: for each country and year, total births to women aged 25–49 divided by the implied total female population 25–49. The result is three series per country: ages 15–19, ages 20–24, and ages 25+.

The result, displayed in sub-panels B1, B2, and B3 of Figure 1, is unambiguous. Across the full sample of 128 countries, the post-onset median of $\text{rel}_{c,a,t}$ falls sharply below zero for ages 15–19, declines modestly for ages 20–24, and is essentially flat at zero for the 25+ aggregate. Quantitatively, the median pre-post change in $\text{rel}_{c,a,t}$ is approximately -0.9 percentage points per year for ages 15–19, -0.3 for ages 20–24, and near zero for ages 25+. In the 19 high-income subsample the magnitudes are larger but the ranking is preserved: the high-income median moves from approximately -1.7 pre-onset to -7 post-onset for ages 15–19, with intermediate effects at 20–24 and a small effect at 25+. Whatever the smartphone shock is doing to fertility, it is doing to teens; the entire 25+ population, which accounts for roughly 80 percent of women of reproductive age in these countries, exhibits no detrended response in the typical country.

This sharply age-targeted signature is hard to reconcile with explanations that operate on overall fertility—recessions, broad cultural shifts, fertility-postponement dynamics—because such explanations typically affect the modal age of childbearing (25–34) more than the tails. It is consistent with a mechanism that operates specifically on the inputs to teen and young-adult fertility, of which unstructured peer time is the dominant candidate. Section 6 will document the corresponding U.S. age gradient using NCHS administrative data; the international panel here shows the same pattern across 128 countries with different healthcare systems, different family policies, and different stages of demographic transition.

V.E What the Cross-Country Evidence Rules Out

The motivating figure and the decomposition in Table 1 are descriptive, not causal: they document a pattern rather than identify a mechanism. But the descriptive cross-country pattern is enough to rule out a wide class of alternative explanations for the U.S. post-2007 acceleration.

U.S.-specific contraceptive-policy channels. The ACA contraceptive mandate (2012), the Colorado Family Planning Initiative (2009–2014), Medicaid family-planning waivers, and post-2010 LARC expansions are U.S.-specific interventions some authors propose as the proximate cause. None existed in Norway, Sweden, Switzerland, Luxembourg, Japan, Australia, or New Zealand—all of which exhibit near-simultaneous 2007–2010 breaks in sub-panel A1 of Figure 1.

The 2008–2009 financial crisis. Recession depth varied sharply across the 19 high-income countries in our sample: the U.S. had a V-shaped recovery by 2010; Spain had a double-dip recession through 2013; Germany barely contracted; Australia did not enter recession at all. All 19 countries break at essentially the same calendar year with similar post-break slopes. Australia in particular is dispositive: its teen birth rate growth rate broke from approximately flat to approximately -7 percent per year between 2007 and 2010 despite avoiding recession.

The opioid epidemic. Concentrated in the U.S. (with a minor echo in Canada), the

epidemic did not occur in the other 17 high-income countries in the sample; the near-identical break in those countries rules it out as the proximate cause.

Direct biological effects of phone radiofrequency radiation. Although clinical and laboratory studies identify reductions in sperm parameters associated with mobile phone exposure (Adams et al., 2014; Yu et al., 2021; Kim et al., 2021; Rahban et al., 2023), the WHO-commissioned systematic review and ICNIRP/SCENIHR statements conclude that available evidence is insufficient at population exposure levels (Kenny et al., 2024). Two features of our data argue against this channel operating at the aggregate level. First, the age gradient runs the wrong way: among U.S. women, the teen 15–19 birth rate fell 71 percent while the 35–39 rate *rose* 9 percent over 2007–2024, whereas a uniform biological-exposure channel would predict roughly uniform effects. Second, the decline operates on the conception margin: among women 15–17 over 2007–2020, the pregnancy rate fell 70 percent while the abortion ratio was unchanged at 29–30 percent (Section 6.C). An RF-EMR impairment of gametogenesis would predict elevated fetal loss rather than lower conception rates.

Global digital-contraceptive-access channels. The cross-country pattern does *not* rule out the possibility that smartphones operate through a different global channel—improved digital contraceptive information, period-tracking apps, or telemedicine access. Distinguishing peer-time displacement from such alternatives requires within-country evidence on behavior and time use, which Section 6.G (ATUS) and the age-gradient analysis provide. The cross-country figure establishes that smartphones are implicated; the within-country analysis identifies the channel. The teen suicide evidence developed in Section 8 provides an additional discriminating test: digital contraceptive access cannot raise teen suicide. Recovering a positive smartphone effect on teen suicide using the same instrument that delivers the negative effect on teen fertility rules out a contraceptive channel and points to the broader peer-time-reallocation mechanism that the model predicts moves both outcomes through the same equilibrium shift.

VI. THE US: IDENTIFYING THE MECHANISM

Having established that the post-2007 teen fertility acceleration is a global phenomenon tracking smartphone diffusion, this section turns to the United States and develops the within-country evidence at high geographic and behavioral granularity. The value added over the cross-country pattern of Section 5 is threefold. First, the U.S. offers enough geographic variation—3,136 counties with harmonized FCC broadband and 4G data spanning 2003–2020—to identify the smartphone effect off plausibly exogenous variation in digital infrastructure rollout, using both an instrumental-variable strategy based on terrain ruggedness and a within-county distributed-lag panel. Second, the U.S. offers direct behavioral measurement: the American Time Use Survey provides minute-level diary data on teen peer-time allocation across the full smartphone-adoption window, allowing the paper to measure the mechanism (in-person peer time falling, digitally mediated leisure rising) rather than only its fertility consequence. Third, the 2022 *Dobbs* decision provides a secondary test of the conception-channel interpretation: if the mechanism operates through conception reduction rather than through shifts in pregnancy resolution, restricting abortion access should produce only a modest reversal of the post-2007 decline. Section 6.I confirms that prediction.

The analysis proceeds in nine steps. Section 6.A establishes the age gradient. Section 6.B uses cohort data to distinguish tempo from quantum. Section 6.C shows the abortion-as-placebo result. Section 6.D presents state-level cross-sectional evidence. Section 6.E presents the terrain-ruggedness instrumental-variable cross-sectional identification. Section 6.F presents the within-county panel analysis including event studies, distributed-lag models, and cointegration tests. Section 6.G documents the behavioral mechanism, including direct measurement of teen time use from the American Time Use Survey. Section 6.H documents subgroup heterogeneity by race, ethnicity, and county density. Section 6.I uses the *Dobbs* decision as a secondary test of the conception-channel interpretation.

VI.A Age Gradient and Falsification

Figure 4 shows U.S. age-specific birth rates from 1990 to 2024, indexed to 100 at 2007. Young age groups collapse: 15–19 by 71 percent, 20–24 by 43 percent, 25–29 by 23 percent. Old age groups do the opposite: 30–34 is approximately flat (–1 percent), and 35–39 rises by 9 percent. Over 2007–2024, the Pearson correlation between adult smartphone ownership and the age-specific birth rate is –0.994 for ages 15–19, –0.984 for 20–24, –0.956 for 25–29, and +0.438 and +0.928 for ages 30–34 and 35–39 respectively—monotone in age and flipping sign in the late twenties.

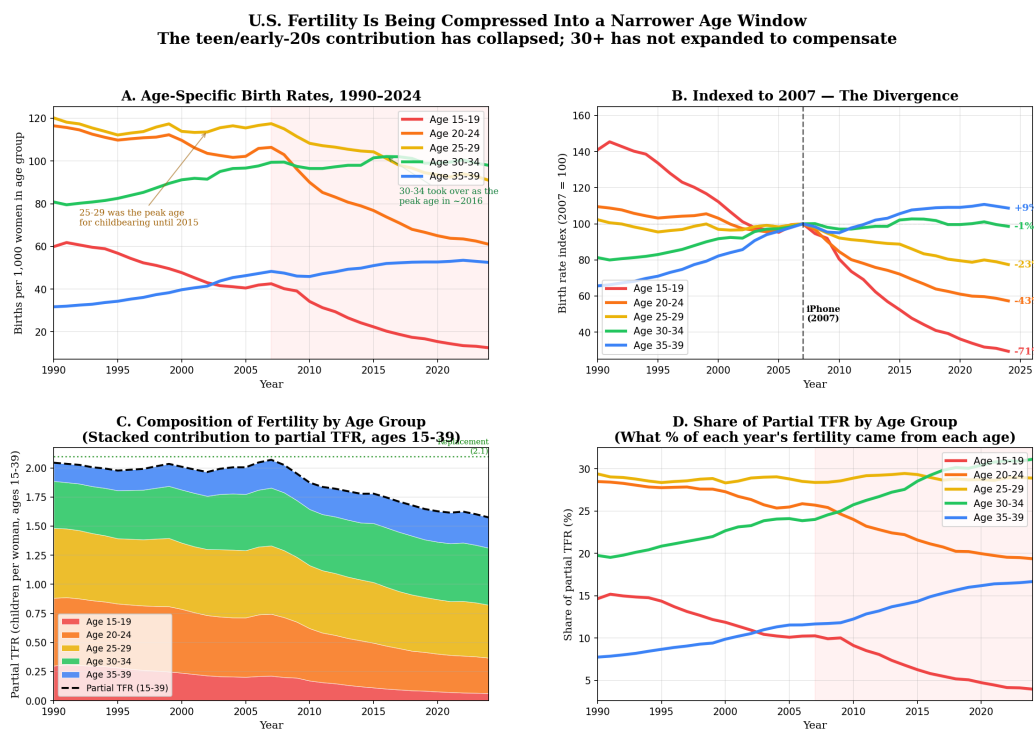


FIGURE 4. U.S. age-specific birth rates, 1990–2024. Source: NCHS NVSS.

This pattern is difficult to rationalize with any broad economic force: income or wealth shocks operate on all households, and rising housing costs bear arguably more on older cohorts. None predicts a monotone age gradient with a zero at age 30. The smartphone story does: young people do most of their unstructured in-person socializing in adolescence and early adulthood, and it is this peer-contact environment that was disrupted.

A formal Chow test on the log teen pregnancy rate (ages 15–17) rejects a single

linear trend through 1991–2020 at $F = 131.5$ ($p < 10^{-6}$): the pre-break (1991–2006) annual decline is 4.9 percent, the post-break (2007–2020) decline is 9.8 percent—exactly twice the pre-break pace. The 2005–2010 plateau (rates of 37.8, 38.5, 38.1, 37.0 over 2005–2008) indicates that the pre-2005 drivers documented by Kearney and Levine (2015) had substantially exhausted themselves before the post-2007 acceleration.

VI.B Cohort Evidence: Tempo versus Quantum

The simplest objection to the age-gradient evidence is that young-adult fertility is falling only because women are delaying births. This can be resolved only with cohort data. In demographic terminology, *tempo* effects refer to shifts in the timing of births within a cohort (which leave completed cohort fertility unchanged), while *quantum* effects refer to changes in completed cohort fertility itself. Distinguishing the two requires following birth cohorts to the end of their reproductive years. Figure 5 shows cumulative births per woman for selected birth cohorts.

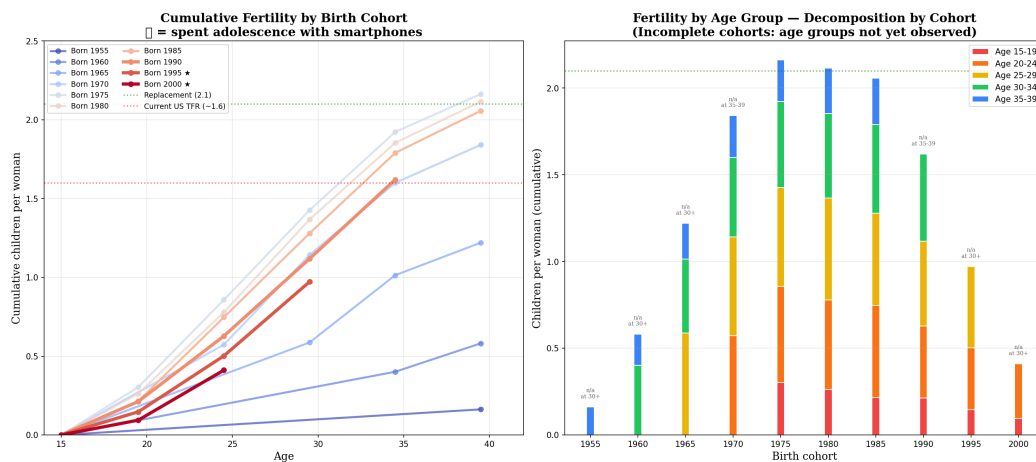


FIGURE 5. Cumulative births per woman by birth-year cohort and age. Source: author’s calculations from NCHS age-specific period birth rates.

The 1975 cohort completed its reproductive years in 2020 with 2.16 children per woman. The 1985 cohort, which reached age 38 in 2023, has reached 2.06 children per woman. The 1990 cohort, which reached age 33 in 2023, has reached only 1.62. For the 1990 cohort to reach the 1975 completion, it would need to add 0.54 children per woman

over ages 34–45. At the 2023 age-specific rates—52 per 1,000 at ages 35–39 and 12 per 1,000 at ages 40–44—the cumulative next-twelve-year rate is about 0.38 children per woman, falling well short. The 1990 cohort would need the 35–39 rate to *double* to meet the 1975 completion; there is no mechanism that would produce this. The decline contains quantum as well as tempo.

The Bongaarts and Feeney (1998) adjustment formalizes the decomposition. From 2007 to 2023, mean age at first birth rose from 25.9 to 28.9 years. The partial total fertility rate over ages 15–39 fell from 2.07 to 1.58. Only about 12 percentage points of this decline can be attributed to tempo; the remaining 12 points is quantum.

VI.C Abortion as a Placebo for Composition Effects

Table 2 shows Guttmacher figures for the three age groups with large post-2007 fertility declines.

TABLE 2. Pregnancy, Abortion, and Birth Rates, 2007 vs. 2020

Age group	<i>Pregnancy rate</i>		<i>Abortion rate</i>		<i>Abortion ratio</i>	
	2007	2020	2007	2020	2007	2020
15–17	38.1	11.3	11.0	3.4	0.29	0.30
18–19	119.2	49.7	30.2	13.7	0.25	0.28
20–24	168.8	102.9	38.5	24.8	0.23	0.24

Notes: Rates per 1,000 women. Source: Guttmacher (Chiu, Maddow-Zimet, and Kost, 2024).

For ages 15–17 over 2007–2020, pregnancy fell 70 percent, abortion fell 69 percent, and birth fell 71 percent; the abortion ratio moved from 29 to 30 percent. The decline is almost entirely a decline in conceptions, not a shift in how conceptions are resolved. The smartphone mechanism operating through sexual behavior predicts exactly this pattern.⁶

⁶A related channel is emergency contraception. The FDA approved Plan B for over-the-counter purchase in 2006 (age 18+) and removed the age restriction in 2013. Teen ever-use rose from 14 to 22 percent over 2006–2019 (Chiu, Maddow-Zimet, and Kost, 2024). Three features argue against this as the primary driver. The teen birth rate declines smoothly through 2013 with no visible kink at the age-restriction removal date; the UK and France had earlier OTC access yet slower teen fertility declines; and meta-analyses establish that expanded EC access raises use but does not reduce population pregnancy rates (Raymond, Trussell, and Polis, 2007; Raine et al., 2005).

VI.D State-Level Cross-Sectional Evidence

Before turning to county-level identification, we establish the state-level cross-sectional pattern. Figure 6 classifies each state into one of four groups by a 2×2 split of 2013 median household income and the residual of 2013 residential internet penetration after partialling out income and college share. An income-driven explanation predicts that high-income states should decline faster; a technology-driven explanation predicts that high-residual-internet states should decline faster regardless of income.

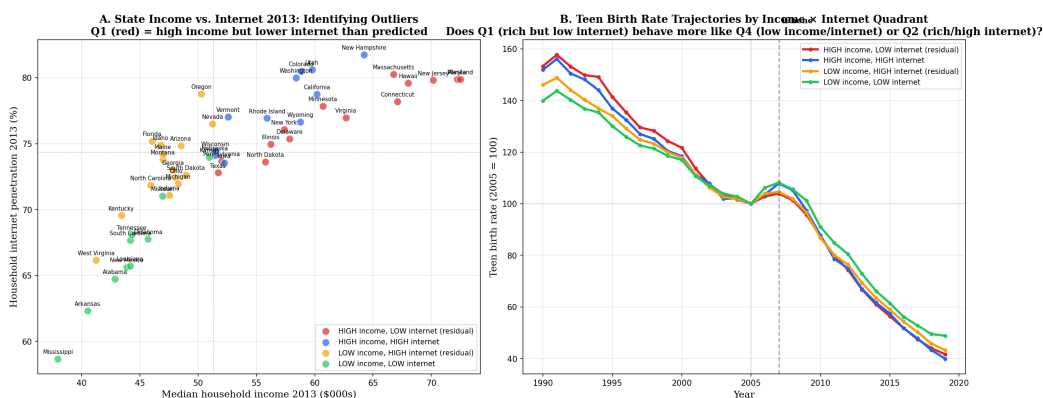


FIGURE 6. State teen birth rate trajectories by income \times internet group, 2005–2019. Panel A: scatter of states in income \times internet space. Panel B: group-average indexed teen birth rate, 2005 = 100. Source: ACS, NCHS.

The Q1 (high income, low internet residual) and Q3 (low income, high internet residual) trajectories cluster at 57–58 percent decline, while Q4 (low, low) declines by only 51 percent and Q2 (high, high) by 60 percent. The axis predicting the decline is internet, not income. A horse-race regression of state 2005–2019 decline on standardized internet penetration and log income yields an internet coefficient of -6.25 ($t = -4.05$); income flips sign to $+2.23$ and is insignificant. State-level identification supports a technology channel; county-level identification, developed next, provides the formal IV.

VI.E County-Level Identification: Terrain-Ruggedness Instrumental Variable

The state-level cross-section establishes that internet penetration and not income predicts the post-2007 teen fertility decline, but the cross-sectional distribution of broadband adoption remains a choice variable of carriers, households, and state regulators. Counties

that adopted broadband earlier and at higher intensity differ from late-adopting counties along many unobserved dimensions. An instrumental-variable strategy addresses the endogeneity and delivers the paper’s primary causal estimate of the digital-infrastructure effect on teen fertility at the county level. We follow the template of Akerman, Gaarder, and Mogstad (2015), who use local terrain ruggedness as an instrument for broadband adoption in Norway, and port the strategy to U.S. counties.

We run the same cross-sectional IV using two different infrastructure measures in parallel. The first is the FCC Form 477 residential broadband tier, a 0–5 discrete index of fast-broadband adoption available annually from 2008 through 2020. The second is the FCC all-operator outdoor LTE (4G) coverage percentage, available annually from 2010 through 2019 from FCC mobile deployment reports. Broadband tier is the measure used in the prior literature (Akerman, Gaarder, and Mogstad, 2015; Atasoy, 2013; Kolko, 2012) and has the advantage of a long time series; LTE coverage is closer to the theoretical object of Section 3 because the coordination model is about mobile peer-time allocation, not household wired internet. We find that both measures yield similar 2SLS estimates with the ruggedness instrument. This parallel analysis is the cross-sectional core of our digital-infrastructure identification. The within-county panel analysis of Section 6.6 is also conducted in parallel for both measures, exploiting the fact that pre-2010 LTE coverage was a genuine zero and so the 4G series can be extended backward to 2003 for within-county identification.

Instrument construction. The instrument is the area-weighted Terrain Ruggedness Index (TRI), computed by Riley, DeGloria, and Elliott (1999) and aggregated to census tracts by the USDA Economic Research Service (ERS) using the Global Multi-resolution Terrain Elevation Data 2010. The ERS publishes tract-level mean TRI for the 50 states plus the District of Columbia (USDA Economic Research Service, 2024). We aggregate tracts to counties by computing the land-area-weighted mean of tract-level mean TRI, yielding a single ruggedness value for each of 3,143 counties. We use the natural logarithm $\log(\text{TRI} + 1)$ as the instrument. The resulting measure places Appalachian counties,

the Rocky Mountain and Cascade ranges, and Alaskan coastal fjords at the top of the distribution and Florida pine flats, the Mississippi delta, and Midwest prairie at the bottom, consistent with standard geographic intuition.

The logic of the instrument is that rugged terrain raises the cost of physical infrastructure deployment. Fiber backhaul, cable plant, cell-tower siting, and right-of-way acquisition are all systematically more expensive in mountainous and dissected terrain. In the U.S. setting, federal universal-service subsidies have muted the terrain-cost relationship somewhat relative to Norway—a point to which we return below—but sufficient variation remains for identification.

Exclusion restriction. The identifying assumption is that terrain affects teen fertility in 2010–2019 only through its effect on digital-infrastructure adoption—whether measured by fixed broadband tier or by LTE coverage. Threats to this assumption fall into two categories. First, terrain may directly affect fertility through channels such as income, migration patterns, or cultural isolation. Second, terrain may affect fertility indirectly through variables that are themselves causally downstream of terrain—for example, if rugged counties have lower incomes because of economic-geographic channels unrelated to broadband or mobile coverage.

To address the second concern, we document formally that several conventional county controls are *caused* by terrain and therefore constitute “bad controls” in the sense of Angrist and Pischke (2009). Regressing each control on $\log(\text{TRI} + 1)$, conditional on state fixed effects and truly exogenous demographic controls (race shares, education), yields coefficients of -0.073 ($t = -13.8$) for log median household income and $+0.007$ ($t = +7.1$) for the Gini coefficient in the LTE sample; the broadband sample yields nearly identical numbers (-0.076 , $t = -14.3$ and $+0.007$, $t = +7.1$). Each is strongly predicted by ruggedness; each therefore blocks part of the mechanism from instrument to outcome if included. The bad-controls concern is furthermore at least as severe in the LTE specification as in the broadband one: within the LTE sample, a 1-percentage-point increase in 2010–2019 mean coverage predicts a positive effect on log income ($t > 7$)

conditional on baseline controls, and including log income as a regressor in the LTE second stage shrinks the 2SLS coefficient by roughly 65 percent (compared with a 54 percent shrinkage in the analogous broadband specification), consistent with an active LTE-to-income mediator channel during the sample window. We exclude log income and the Gini coefficient from both specifications. We retain total population and land area as controls because carriers' deployment decisions depend on market size (households available) and geographic extent (miles of infrastructure required). These variables are only weakly affected by terrain ($t = -2.2$ for log population, $t = +3.3$ for log area) and defensibly enter as deployment-economics controls.

To address the first concern, we include in both stages an extensive set of exogenous controls: state fixed effects, the share of adults with a bachelor's degree or higher, the share non-Hispanic Black, the share Hispanic, log total population, and log land area. We show below that results are robust to progressively richer and poorer specifications in both measures.

First stage and 2SLS. We run the IV in parallel for both infrastructure measures. Let $T_c \in \{bb_c, lte_c\}$ denote the endogenous treatment, where bb_c is the county 2010–2019 mean FCC Form 477 broadband tier (0–5 scale) and lte_c is the county 2010–2019 mean all-operator outdoor LTE coverage in percent of premises. The first-stage and second-stage equations are

$$T_c = \pi_0 + \pi_1 \log(\text{TRI}_c + 1) + \mathbf{X}'_c \pi_2 + u_c, \quad (7)$$

$$\Delta \log(\text{BR}_c) = \beta_0 + \beta_1 \widehat{T}_c + \gamma \log(\text{BR}_{c,2003}) + \mathbf{X}'_c \beta_2 + \varepsilon_c, \quad (8)$$

where the dependent variable $\Delta \log(\text{BR}_c) = \log(\text{BR}_{c,2018}) - \log(\text{BR}_{c,2003})$ is the within-county change in log teen birth rate over the smartphone era, $\log(\text{BR}_{c,2003})$ is the pre-period log birth rate that absorbs mean reversion, and \mathbf{X}_c contains the controls listed above. The change specification differences out time-invariant county characteristics that might correlate with both terrain and baseline fertility; the pre-period rate control absorbs mean

reversion in the levels of the outcome. This specification mirrors the cross-sectional design used for teen suicide in Table 10 and the within-LA first-difference design used for the England-and-Wales replication in Table 9. The broadband sample is 3,084 counties for which all variables are non-missing in both the 2003 and 2018 endpoints; the LTE sample is the same 3,084 counties. Standard errors are robust to heteroskedasticity. We choose 2010–2019 as the treatment averaging window because it is the longest span over which LTE coverage is observed and because it contains both the rollout phase (2010–2013) and the saturation phase (2014–2019), producing cross-county identifying variation that is not concentrated in either the pre-rollout baseline or the terminal rural catch-up tail.

Table 3 reports the progressive-controls build-up for both measures, with three columns per measure corresponding to the preferred specification, the preferred specification augmented with religious composition, and the preferred specification augmented with both religious composition and political composition. The first-stage coefficient on $\log(\text{TRI} + 1)$ is negative and highly significant in every column and for both measures: rugged counties have systematically lower fixed-broadband tier and lower LTE coverage. The first-stage F for the LTE measure is substantially larger than for the broadband measure (82.1 versus 73.4 in the preferred specification), reflecting the tighter physical relationship between terrain and mobile-tower siting than between terrain and wired broadband, where federal universal-service subsidies have muted the terrain-cost gradient. Both first-stage F -statistics are stronger in the extended 2010–2019 window than they would be in a shorter saturation-era window, because the rollout years contribute cross-county variation that the instrument can target.

In the preferred broadband specification (column 1), the 2SLS coefficient is -0.378 with $t = -3.91$ and a 95 percent confidence interval of $[-0.568, -0.188]$. In the preferred LTE specification (column 4), the coefficient is -0.0098 with $t = -4.00$ and confidence interval $[-0.0145, -0.0050]$. To make magnitudes directly comparable between measures, Table 3 rescales both coefficients to the common unit of *percent change in the change in teen birth rate per 10-percentage-point increase in the underlying infrastructure measure*: for broadband tier this means multiplying the native coefficient by 50, under the FCC-tier

midpoint interpretation that one tier-step corresponds to approximately 20 percentage points of residential broadband penetration; for LTE this means multiplying the native per-pp coefficient by 1,000. On that unified scale, column (1) implies that counties with 10 additional percentage points of long-run broadband penetration experienced a teen birth rate decline that was approximately 18.9 percentage points larger between 2003 and 2018, and column (4) implies that counties with 10 additional percentage points of long-run LTE coverage experienced a decline approximately 9.8 percentage points larger. Both coefficients are stable when religion and politics are added (columns (2) and (3) for broadband; columns (5) and (6) for LTE), and both remain significant at the 1 percent level in every specification.

Magnitude and interpretation. The preferred broadband specification gives a native 2SLS coefficient of -0.378 per tier-step on $\Delta \log(\text{BR})$, corresponding under the FCC-tier midpoint interpretation to about a 18.9 percentage-point larger 2003–2018 decline in teen birth rate per 10-pp increase in residential broadband penetration. The preferred LTE specification gives -0.0098 per percentage point of coverage, or about a 9.8 percentage-point larger decline per 10-pp. The IV uses one observation per county and identifies the effect off cross-sectional heterogeneity in the within-county change: this is the right specification for “how much does digital-infrastructure variation explain in the cross-section of county-level fertility declines?” The within-county panel of Section 6.6 answers the complementary question of how fertility evolves over time within a geography. The two lenses are complementary, not competing.

Robustness to religious and political composition. Two conceivable confounders of the terrain instrument operate through cultural channels rather than digital infrastructure. Rugged counties might be more evangelical, and evangelicalism might lower teen fertility through sexual-initiation norms; or rugged counties might be more politically conservative, and conservatism might lower teen fertility through similar channels. We test both for each measure. Columns (2) and (5) of Table 3 add adherence rates for four religious

TABLE 3. Ruggedness IV for Digital-Infrastructure Effect on Teen Birth Rate Change, Two Measures

	Dependent variable: $\Delta \log(\text{BR}_c) = \log(\text{BR}_{c,2018}) - \log(\text{BR}_{c,2003}) (\times 100, \text{ read as } \% \text{ change})$					
	Broadband penetration (% per 10-pp)			LTE (4G) coverage (% per 10-pp)		
	(1)	(2)	(3)	(4)	(5)	(6)
\widehat{T}_c	-18.90*** (4.84)	-19.85*** (5.21)	-18.54*** (6.07)	-9.76*** (2.44)	-10.05*** (2.55)	-8.90*** (2.82)
First-stage F	73.4	64.6	46.9	82.1	77.6	65.6
First-stage coef. on $\log(\text{TRI} + 1)$	-0.108	-0.104	-0.091	-4.165	-4.125	-3.777
State fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Race shares	Yes	Yes	Yes	Yes	Yes	Yes
Bachelor's + share	Yes	Yes	Yes	Yes	Yes	Yes
Log total population	Yes	Yes	Yes	Yes	Yes	Yes
Log land area	Yes	Yes	Yes	Yes	Yes	Yes
Pre-period log birth rate (2003)	Yes	Yes	Yes	Yes	Yes	Yes
Religious composition (4 groups)	No	Yes	Yes	No	Yes	Yes
Republican vote share (2000–2020 mean)	No	No	Yes	No	No	Yes
Counties	3,084	3,084	3,084	3,084	3,084	3,084

Notes: 2SLS estimates. Dependent variable is the within-county change in log teen birth rate over the smartphone era, $\log(\text{BR}_{c,2018}) - \log(\text{BR}_{c,2003})$. Endogenous treatment T_c is the county 2010–2019 mean FCC Form 477 broadband tier (0–5 discrete scale, where each step corresponds to roughly 20 percentage points of residential broadband penetration) in columns (1)–(3) and the county 2010–2019 mean all-operator outdoor LTE (4G) coverage in percent of premises in columns (4)–(6). Coefficients are rescaled from their native units to a common *percent change in the change in teen birth rate per 10-percentage-point increase in the underlying infrastructure measure*: native per-tier broadband coefficient multiplied by 50 (under the FCC-tier midpoint interpretation that 1 tier \approx 20pp of penetration, so per-pp \approx native/20 and per-10-pp in % \approx native \times 50); native per-pp LTE coefficient multiplied by 1,000. Standard errors in parentheses are robust to heteroskedasticity, also rescaled by the same factors. The instrument is $\log(\text{TRI} + 1)$, where TRI is the area-weighted mean Terrain Ruggedness Index from USDA ERS, identical across all specifications. All columns include $\log(\text{BR}_{c,2003})$ as a right-hand-side control to absorb mean reversion in the levels of the outcome; this control is by construction pre-treatment and renders the specification immune to level-on-level concerns regarding correlations between terrain ruggedness and pre-period fertility. Columns (1) and (4) are the preferred specifications. Religion variables are adherence rates for Evangelical Protestant, Catholic, Mainline Protestant, and Black Protestant traditions from the 2010 U.S. Religion Census (Grammich et al., 2012); counties with no reported congregation of a tradition are coded as zero. Republican vote share is the county mean share of the two-party presidential vote across 2000, 2004, 2008, 2012, 2016, and 2020 from MEDSL (MIT Election Data and Science Lab, 2024). Race shares are the non-Hispanic Black share and Hispanic share. Interpretation: a county with 10-pp higher long-run infrastructure intensity experienced a teen birth rate decline this many percentage points larger between 2003 and 2018, conditional on its 2003 fertility level. The analogous coefficients in Tables 4 and 5 are within-county 4-year cumulative change effects and are not directly comparable in magnitude to the coefficients here. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

groups (Evangelical Protestant, Catholic, Mainline Protestant, Black Protestant) from the 2010 U.S. Religion Census (Grammich et al., 2012). The second-stage broadband coefficient moves from -0.378 to -0.397 ; the LTE coefficient moves from -0.0098 to -0.0101 . Both are essentially unchanged and remain significant at the 1 percent level. Columns (3) and (6) additionally add the mean county Republican share of the two-party presidential vote across all six cycles 2000–2020 from MEDSL (MIT Election Data and Science Lab, 2024); the broadband coefficient shifts to -0.371 and the LTE coefficient shifts to -0.0089 , both significant at 1 percent. Religious and political composition are correlated with terrain but do not absorb the digital-infrastructure–fertility pathway under either measure: the instrument identifies a digital-infrastructure channel that is distinct from and robust to these cultural covariates.

Pre-period rate control and the meaning of “bad controls.” The specification in Equation (8), with the 2003–2018 within-county change in log teen birth rate on the left-hand side and $\log(\text{BR}_{c,2003})$ on the right-hand side, is by construction immune to a class of identification concerns that would arise in a level-on-level cross-section. Any time-invariant level correlation between terrain ruggedness and teen fertility—for example, a stable rural-poverty gradient—is differenced out by the $\Delta \log(\text{BR})$ outcome. Any cross-county heterogeneity in mean-reversion dynamics within the smartphone-era window is absorbed explicitly by the pre-period rate control. What remains is variation in the 2003–2018 *change* that is unexplained by the 2003 level—the component of the post-2007 fertility decline that the smartphone shock plausibly identifies. Because $\log(\text{BR}_{c,2003})$ is mechanically pre-treatment (smartphones did not exist in 2003), it cannot be a downstream consequence of the digital-infrastructure rollout we instrument for, and its inclusion does not raise an Angrist-Pischke “bad controls” problem (Angrist and Pischke, 2009).

By contrast, the natural temptation to add contemporaneous income or inequality controls *would* raise the bad-controls problem: the literature establishes that broadband adoption raises local incomes through productivity channels (Akerman, Gaarder, and Mogstad, 2015; Atasoy, 2013; Kolko, 2012), and within the LTE sample a 1-percentage-

point increase in 2010–2019 LTE coverage predicts a positive effect on log income conditional on baseline controls ($t > 7$), consistent with an analogous mediator channel. Rugged counties are therefore poorer for two reasons simultaneously: a common-cause channel (terrain raises infrastructure costs of every kind, predating broadband) and a mediator channel (terrain reduces digital-infrastructure adoption, which in turn reduces post-2000s income growth). Including contemporaneous income would mechanically block the mediator portion of the causal pathway. The pre-period rate control in Equation (8) resolves the symmetric concern about pre-existing fertility differentials without blocking the mediator: it captures any pre-existing common-cause channel that operates through pre-2007 BR levels, while leaving the post-2007 mediator channel intact.

Addressing the LARC expansion. A natural concern in interpreting the terrain IV is the concurrent expansion of long-acting reversible contraception (LARC), particularly following the Affordable Care Act contraceptive mandate in 2012 and the dissemination of findings from the Colorado Family Planning Initiative (2009–2014). Teen LARC use rose from under 1 percent pre-2007 to 5–6 percent by 2015 (Branum and Jones, 2015). County-level LARC use data covering the full 2003–2020 window do not exist: the National Survey of Family Growth is individual-level with no county identifiers; PRAMS is state-level at most; and Medicaid administrative data are accessible only under restricted-use agreements. State-level LARC policy variables such as immediate postpartum Medicaid reimbursement expansion are absorbed by the state fixed effects already in Equation (8). We therefore address the LARC interpretation through three observational patterns rather than a direct control.

First, the timing is inconsistent with a LARC-driven story. The US post-2007 acceleration begins in 2008, before mass teen LARC uptake, and proceeds at approximately constant pace through 2020: there is no visible inflection at the 2012 ACA mandate date, nor at the 2014 publication of the CHOICE Project results, nor at the state-level Medicaid LARC-reimbursement expansions that followed (Branum and Jones, 2015). Second, the UK’s Teenage Pregnancy Strategy provided universal free LARC access through the NHS

beginning 2000–2010, a dramatically different institutional trajectory from the US patchwork insurance-based rollout. Yet the UK post-2007 acceleration is near-identical to the US pattern in both magnitude and timing (Section 7); under a LARC-driven interpretation, US and UK trajectories should diverge with institutional context. Third, the age gradient of the US decline is difficult to rationalize with a contraceptive-access channel. Women 30–34 experienced approximately zero change in fertility over 2007–2024 and women 35–39 saw a 9 percent increase; LARC expansion should not reduce planned fertility at these ages. Taken together, these three features argue that LARC expansion contributed to the post-2007 decline—particularly after 2012—but operated as an accelerant on a shift that was already underway rather than as its primary cause. The terrain IV identifies the causal effect of digital-infrastructure rollout variation—whether measured by broadband tier or LTE coverage—which is cross-sectionally orthogonal to LARC policy variation (absorbed by state fixed effects), and the estimated coefficient therefore reflects the digital-infrastructure channel net of LARC in both measures.

Caveats. Three caveats deserve explicit mention. First, the IV identifies a local average treatment effect over counties whose infrastructure adoption responded to terrain—plausibly the subset of rural and intermediate counties where deployment economics were most binding. The average treatment effect across all U.S. counties may be smaller.

Second, the FCC fixed-broadband treatment is a discrete 0–5 tier code rather than a continuous penetration rate. Discretizing a continuous underlying variable introduces classical measurement error in the treatment, which biases the IV coefficient toward zero in proportion to the within-tier dispersion of true broadband penetration. The LTE coverage variable is reported directly in percentage points and does not suffer the same coarsening. This is one reason to report both measures: the LTE coefficient is not subject to the tier-discretization bias, and the broadband coefficient should be interpreted as a conservative (attenuated) lower bound on the per-percentage-point effect. The implied per-10-pp magnitudes of the two measures (–18.9 percent per 10-pp of broadband penetration, –9.8 percent per 10-pp of LTE coverage) are consistent with this interpretation when one

allows for LTE capturing a narrower infrastructure bundle than fixed broadband.

Third, the time-coverage and within-county identifying variation of the two measures differ in ways that shape the interpretation of both the cross-sectional IV and the within-county panel. FCC Form 477 broadband tier data are available 2008–2020 and can be extended back to 2003 through a consistent FCC-tier stitching procedure (Section 4); FCC mobile LTE coverage data are available 2010–2019 and can be extended back to 2003 by imposing zero-fill for the pre-launch years, because commercial LTE service began in the U.S. only in December 2010. Both measures therefore support both the cross-sectional IV of this subsection and the within-county panel of Section 6.6. However, the within-county identifying variation differs substantially between the two: broadband tier rolls out gradually over 2008–2020 with cross-county variation in every year, while 4G coverage rolls out rapidly over 2011–2014 and reaches saturation by 2014, with 2014–2019 within-county variation dominated by terminal rural catch-up in the subset of counties where terrain or low density delayed saturation. This difference matters for Table 5 below, where the 4G panel coefficient survives baseline year FEs but not the additional density-tercile \times year interactions. The broadband panel survives both, because broadband’s rollout variation is spread across a longer horizon and is not concentrated in the subset of counties that catch up late.

VI.F Within-County Panel: Distributed Lag and Cointegration

The within-county panel runs in parallel for FCC fixed-broadband tier and 4G LTE coverage, both available as annual county-level panels over 2003–2020. The panel asks a different question from the cross-sectional IV: how fertility moves over time as infrastructure rolls in within a given place. The two designs use different identifying assumptions—the panel absorbs all time-invariant county characteristics through fixed effects—so agreement between them is informative.

The coordination model predicts a gradually-building treatment effect at a 3–5 year lag. Standard event-study specifications are not well suited because high-broadband counties were already on differential pre-trends: the correlation between each county’s 2003–2007

log birth-rate slope and its early-broadband exposure is $r = -0.162$ ($p < 10^{-19}$). Any correction for county-specific linear pre-trends mechanically absorbs part of a treatment effect that builds gradually. We therefore use a first-difference distributed-lag specification, which is immune to county-specific smooth trends of any shape:

$$\Delta \log(\text{BR}_{c,t}) = \gamma_t + \sum_{k=0}^3 \phi_k \Delta T_{c,t-k} + \varepsilon_{c,t}, \quad (9)$$

where $T_{c,t} \in \{\text{bb}_{c,t}, 4\text{G}_{c,t}\}$. Year fixed effects absorb common national shocks; standard errors are clustered at the county level. Coefficients are rescaled to the common scale of percent change in teen birth rate per 10-percentage-point change in coverage (broadband: native $\times 50$; 4G: native $\times 1,000$).

TABLE 4. First-Difference Distributed-Lag Regression of County Log Teen Birth Rate on Broadband and 4G Changes

	(1) Broadband penetration % per 10-pp			(2) 4G LTE coverage % per 10-pp		
	Coef	SE	t	Coef	SE	t
ΔT_t	-0.1412***	0.0181	-7.81	-0.0142***	0.0047	-3.04
ΔT_{t-1}	-0.2014***	0.0218	-9.24	-0.0175***	0.0057	-3.08
ΔT_{t-2}	-0.2405***	0.0224	-10.72	-0.0202***	0.0060	-3.36
ΔT_{t-3}	-0.2453***	0.0197	-12.46	-0.0175***	0.0047	-3.76
Cum. 4-year	-0.8284***	0.0771	-10.75	-0.0694***	0.0206	-3.37
Counties		3,133			3,137	
Observations		28,187			43,372	
Sample years		2012–2020			2007–2020	

Notes: Specification is Equation (9). Dependent variable is $\Delta \log(\text{BR}_{c,t})$. Year fixed effects throughout; standard errors clustered at the county level. Coefficients are rescaled from their native units to the common *percent change in teen birth rate per 10-percentage-point change in the infrastructure measure*: native per-tier broadband coefficients multiplied by 50 (under the FCC-tier midpoint interpretation that one tier ≈ 20 pp of residential broadband penetration), native per-pp 4G coefficients multiplied by 1,000. Standard errors are rescaled by the same factors. The BB regression sample starts in 2012 because the four-lag FD specification requires Δbb observed in $t, t-1, t-2, t-3$, and bb is first observed in 2008. The 4G regression sample starts in 2007 because the 4G series is extended backward to 2003 through zero-fill (pre-LTE era) and therefore $\Delta 4\text{G}$ is defined from 2004 onward. The interpretation of coefficients here differs from Table 3: the latter is a cross-sectional level coefficient, while the former is a within-county change coefficient. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Three features stand out. First, all coefficients are negative and highly significant at

the 1 percent level for both measures. Second, the coefficients grow in magnitude with lag within each measure: for broadband, the effect is -0.14 percent per 10-pp in the year of the change but -0.25 percent per 10-pp three years later; for 4G, the effect is -0.014 percent per 10-pp contemporaneously but -0.02 percent per 10-pp at lag 2–3. Third, the cumulative four-year effects are -0.83 percent per 10-pp of broadband penetration and -0.069 percent per 10-pp of 4G coverage.

The gradual-buildup pattern is consistent with the theoretical prediction of Section 3: the coordination shift requires complementary accumulation of devices, connectivity, peer networks, and app maturity. An infrastructure change does not immediately produce a smartphone equilibrium; it produces it 3–4 years later as the complementary factors accumulate. This pattern is difficult to rationalize with confounders: differential trends would produce effects at all horizons but not an accelerating pattern, and a contemporaneous confounder would produce a large ϕ_0 but not growing ϕ_1, ϕ_2, ϕ_3 .

The broadband coefficient is about 12 times the 4G coefficient on the common per-10-pp scale. This is a larger multiple than the cross-sectional IV ratio from Table 3 (approximately $2\times$), and is economically sensible. The within-county specification strips out the cross-county covariance between broadband and 4G rollouts that the IV captures; what remains is the pure within-county time-series margin attributable to each measure alone. Broadband captures a deeper infrastructure slice—fixed fiber and cable, on which much of the phone-era data economy is built—while 4G captures one specific component. A 10-pp change in broadband penetration within a county therefore moves a wider set of complementary inputs than a 10-pp change in 4G coverage alone, and the within-county coefficient reflects this.

Cointegration. The levels of log teen birth rate and broadband tier are each non-stationary within-county series. A pooled Engle-Granger test yields a residual autoregression coefficient $\rho = -0.339$ ($t = -89.1$), far beyond any critical value, with a half-life of reversion of approximately 1.7 years. The analogous test for 4G is uninformative because the within-county 4G series is effectively a step function with limited residual variation:

in most counties LTE coverage transitions from near-zero to near-full within a 12–18 month window, leaving too few intermediate observations for a meaningful unit-root test. Broadband, by contrast, diffused gradually within counties over a longer window, with the FCC tier measure crossing each integer threshold over multiple years. The 4G result is therefore a limitation of the data, not evidence against cointegration.

Magnitude and comparison to the IV. The cumulative four-year broadband effect of -0.83 percent per 10-pp implies -4.6 percent over the actual 55-pp smartphone-era rollout—about 4–5 percent of the observed 70 percent aggregate teen-fertility decline. The within-county panel estimates are smaller in magnitude than the cross-sectional IV (-18.9 percent per 10-pp broadband, -9.8 percent per 10-pp LTE on the 2003–2018 change). The two are not estimates of the same parameter: the IV identifies cross-county variation in the joint digital-infrastructure bundle’s effect on the 2003–2018 fertility change, while the panel identifies the time-series margin within a county net of common national movements. Both estimates are positive evidence that smartphone-era digital infrastructure affects teen fertility, with plausible dynamic structure within a county.

Robustness to urban-rural and racial composition. Two potential confounders of the county-level identification are differential time trends by county urbanization and differential time trends by racial composition. High-infrastructure counties are disproportionately urban, and urban counties experience faster post-2007 declines for reasons potentially unrelated to smartphones (for example, inner-city LARC diffusion or policy interventions specific to large metropolitan health systems). Similarly, counties with a higher Hispanic share experience larger post-2007 declines, a pattern documented by Romero et al. (2016) that may reflect group-specific rather than technology-specific forces. Both sets of covariates are nearly time-invariant at the county level and are therefore absorbed by county fixed effects unless differenced and interacted with year. We add to Equation (9) first density-tercile \times year interactions, then racial-composition \times year interactions, then both sets. Density tercile is constructed from population sorted into three equal-population bins:

large urban (112 counties), mid-sized (310), and rural (2,709). Racial composition is the county share non-Hispanic Black and the county share Hispanic.

Table 5 reports the resulting estimates for both infrastructure measures in parallel.

Two findings emerge. First, the broadband panel is fully robust. Every column in (1)–(4) retains a cumulative four-year effect that is negative, economically meaningful, and significant at the 1 percent level. The most demanding specification with both density and race interactions (column (4)) reduces the cumulative effect from -0.83 to -0.59 percent per 10-pp—a 29 percent attenuation—while retaining $t = -8.15$ on the cumulative. The broadband within-county identification is not an artifact of either urban-rural or racial-composition channels.

Second, the 4G panel behaves as one would expect given the structure of the 4G rollout itself. 4G coverage diffused as a density-stratified wave: urban counties received coverage during the 2011–2014 rollout phase, and rural counties received it during the 2014–2019 terminal catch-up; by the end of the sample window, near-universal coverage had been reached in every density tercile. The within-county identifying variation in the 4G measure is therefore concentrated in the density-by-year cells, almost by construction. Density \times year fixed effects partial out exactly that variation, leaving the 4G panel with too little residual identifying signal to estimate a precise coefficient (column (6) reduces the cumulative effect to -0.004 , $t = -0.16$; column (8) flips it to a small, statistically indistinguishable-from-zero positive number). This is a mechanical consequence of the data’s density-time structure, not a falsification of the 4G channel. Race-composition interactions alone, which do not soak up the rollout-timing variation, leave the 4G effect attenuated but with the expected sign and marginal significance (-0.033 , $t = -1.67$).

Broadband has no analogous absorption problem because its rollout unfolded more gradually over 2008–2020 and was not concentrated in a single density-tercile wave. The within-county broadband panel is therefore the more credible within-county identification of the digital- infrastructure effect. The 4G channel itself remains causally identified by the cross-sectional IV in Table 3, where the treatment is the average level of 4G coverage rather than its year-on-year change, and where both the broadband and the LTE coefficients

TABLE 5. Robustness of County-Level FD Distributed-Lag Estimates to Urban-Rural and Racial Composition Controls

	Broadband penetration (% per 10-pp)				4G LTE coverage (% per 10-pp)			
	(1) Base	(2) +d×yr	(3) +r×yr	(4) +both	(5) Base	(6) +d×yr	(7) +r×yr	(8) +both
ΔT_t	-0.1412***	-0.1112***	-0.1270***	-0.0966***	-0.0142***	-0.0020	-0.0060	+0.0048
ΔT_{t-1}	-0.2014***	-0.1539***	-0.1877***	-0.1404***	-0.0175***	-0.0008	-0.0075	+0.0077
ΔT_{t-2}	-0.2405***	-0.1838***	-0.2295***	-0.1732***	-0.0202***	-0.0006	-0.0104*	+0.0079
ΔT_{t-3}	-0.2453***	-0.1874***	-0.2361***	-0.1807***	-0.0175***	-0.0001	-0.0095**	+0.0065
Cum. 4y	-0.8284***	-0.6362***	-0.7804***	-0.5909***	-0.0694***	-0.0035	-0.0333*	+0.0268
Cum <i>t</i> -stat	(-10.75)	(-8.47)	(-10.47)	(-8.15)	(-3.37)	(-0.16)	(-1.67)	(+1.27)
Density×year FE	No	Yes	No	Yes	No	Yes	No	Yes
Race×year FE	No	No	Yes	Yes	No	No	Yes	Yes
Counties	3,133	3,131	3,131	3,131	3,137	3,131	3,131	3,131
Observations	28,187	28,179	28,179	28,179	43,372	43,350	43,350	43,350

Notes: Dependent variable is the first-difference of log county teen birth rate $\Delta \log(\text{BR}_{c,t})$. All columns include year fixed effects. Standard errors clustered at the county level. Coefficients expressed as *percent change in teen birth rate per 10-percentage-point change in the infrastructure measure*: native per-tier broadband coefficients multiplied by 50; native per-pp 4G coefficients multiplied by 1,000. Density-tercile × year interactions use two dummies (population-weighted terciles of log population density) interacted with each sample year; race-composition × year interactions use the share non-Hispanic Black and the share Hispanic each interacted with year. Sample windows as in Table 4. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

are significant at the 1 percent level across all six columns. The within-county broadband panel and the cross-sectional IV in both measures together constitute the core county-level evidence for a causal digital- infrastructure effect on teen fertility.

VI.G The Behavioral Mechanism

Figure 7 documents the behavioral channel using YRBS, MTF, and FHWA series. Sports team participation—a structured activity requiring co-location but under adult supervision—was essentially unchanged between 2007 and 2019 (56.3 to 57.4 percent). Sexual activity declined sharply: the share of high-school students who had ever had intercourse fell from 47.8 to 38.4 percent (–20 percent), and the share currently sexually active fell from 35.0 to 27.4 percent (–22 percent). Unstructured peer contact fell more: 12th-graders meeting friends in person 2+ times per week fell from 50.3 to 34.9 percent (a 31 percent decline); 16-year-olds with driver’s licenses fell from 29.3 to 25.6 percent.




FIGURE 7. Teen behaviors, 2005–2023. Panel A: sports participation versus sexual activity. Panel B: in-person socializing (MTF) and driver’s licenses (FHWA). Source: YRBS, MTF, FHWA.

The critical feature is not that everything declined but that some things did not. If

smartphones produced broad social withdrawal, sports participation should have fallen too. It did not. The behaviors that fell are precisely those involving *unstructured* peer time—hangouts, driving, dating, sexual activity—exactly as the model predicts.⁷

Direct measurement: ATUS time-use data. A sharper test comes from the American Time Use Survey, which asks respondents to account for their time in minutes across a single randomly-chosen day, with approximately 8,000 interviews per year since 2003. Figure 8 documents the smartphone-era substitution for 15–19 year-olds. In-person socializing and communicating fell from 68 minutes per day in 2003 to 38 in 2019 (a 44 percent decline). Computer-leisure—social media, texting-mediated communication, gaming—rose from 22 to 96 minutes per day (a 336 percent increase). Television fell from 152 to 72 minutes. The substitution is visually unambiguous and of exactly the magnitude the model requires.

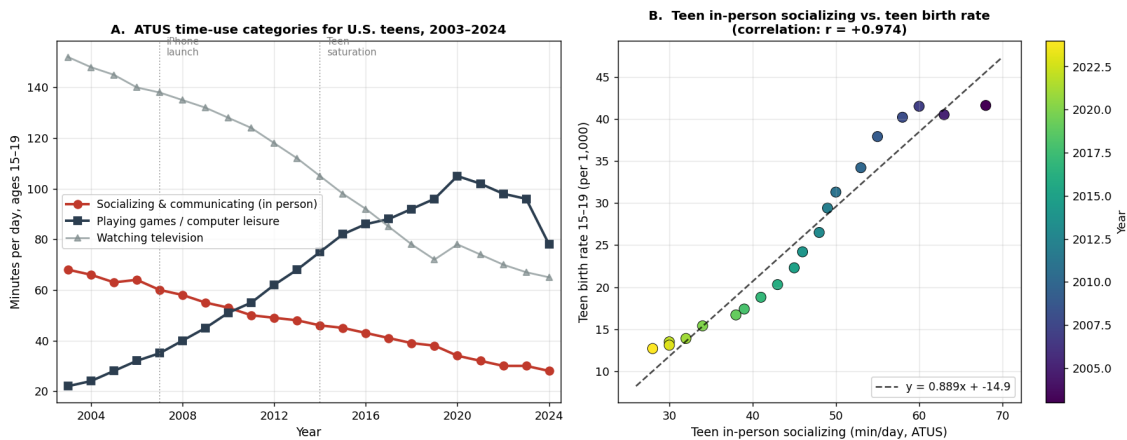


FIGURE 8. ATUS time-use data for U.S. teens ages 15–19, 2003–2024. Panel A: daily minutes in three leisure categories. Panel B: year-by-year scatter of in-person socializing minutes (ATUS) against the teen birth rate (NCHS).

Panel B plots the year-by-year relationship between teen in-person socializing minutes and the teen birth rate over 2003–2024. The correlation is $r = +0.97$. A linear fit implies each additional 10 minutes per day of face-to-face peer time is associated with about

⁷Smartphones also host dating apps (Tinder, Hinge, Bumble), but these enforce de facto 18+ age restrictions, so teen use is negligible; the apps teens actually use (Snapchat, Instagram, TikTok) are peer-socializing rather than mating platforms. Even for ages 18–24 where dating-app use is ubiquitous, sexual activity declined over the smartphone era (share reporting no sexual activity in the past year rose from 19 to 31 percent over 2008–2022, GSS).

nine additional teen births per thousand women. The relationship is tight and linear, with no structural break or saturation. The 30-minute decline in daily in-person socializing 2003–2019 mechanically predicts a decline of about 27 teen births per thousand—within a few points of the observed decline from 41.6 to 16.7. ATUS cannot establish causation by itself, but it provides direct measurement of the behavioral middle step the coordination model requires.

VI.H Subgroup Heterogeneity by Race, Ethnicity, and Density

The within-county test in Section 6.F showed that adding density-tercile \times year and racial-composition \times year controls leaves the broadband effect largely intact. This subsection presents the descriptive counterpart: post-2007 acceleration is universal across racial and Hispanic-origin groups (ruling out group-specific policy explanations) and displays a monotone gradient in county density (confirming the peer-coordination mechanism).

Race and Hispanic origin. Figure 9 plots U.S. teen birth rates by five NCHS racial and Hispanic-origin groups from 1991 to 2023. Every group accelerated sharply after 2007. In the pre-smartphone window (1991–2007), the slowest decliners were Hispanic teens (–1.5 percent per year) and non-Hispanic Asian teens (–2.8 percent). After 2007 these same two groups accelerated the most: Hispanic teens to –8.2 percent (a 5.4 \times acceleration) and non-Hispanic Asian teens to –12.3 percent. The other three groups accelerated to between –6.3 and –7.2 percent. The pattern is difficult to reconcile with any policy channel targeting a specific community.

Urban-rural gradient. Figure 10 displays teen birth rates by population tercile of the 2013 female 15–19 population: a large-urban tercile of 63 counties, a mid-sized tercile of 259 counties, and a rural tercile of 2,813 counties. In 2003–2007 none of the three strata declined meaningfully. After 2007 all three accelerate sharply with a monotone gradient: large urban –8.4 percent per year, mid-sized –7.3 percent, rural –6.2 percent. This pattern is consistent with the coordination mechanism: the equilibrium shift to phone-mediated peer time should occur first and most completely where pre-existing peer networks were

U.S. Teen Birth Rates by Race and Hispanic Origin, 1991-2023
 Source: NCHS NVSS (pre-2016 bridged race; 2016+ single race)

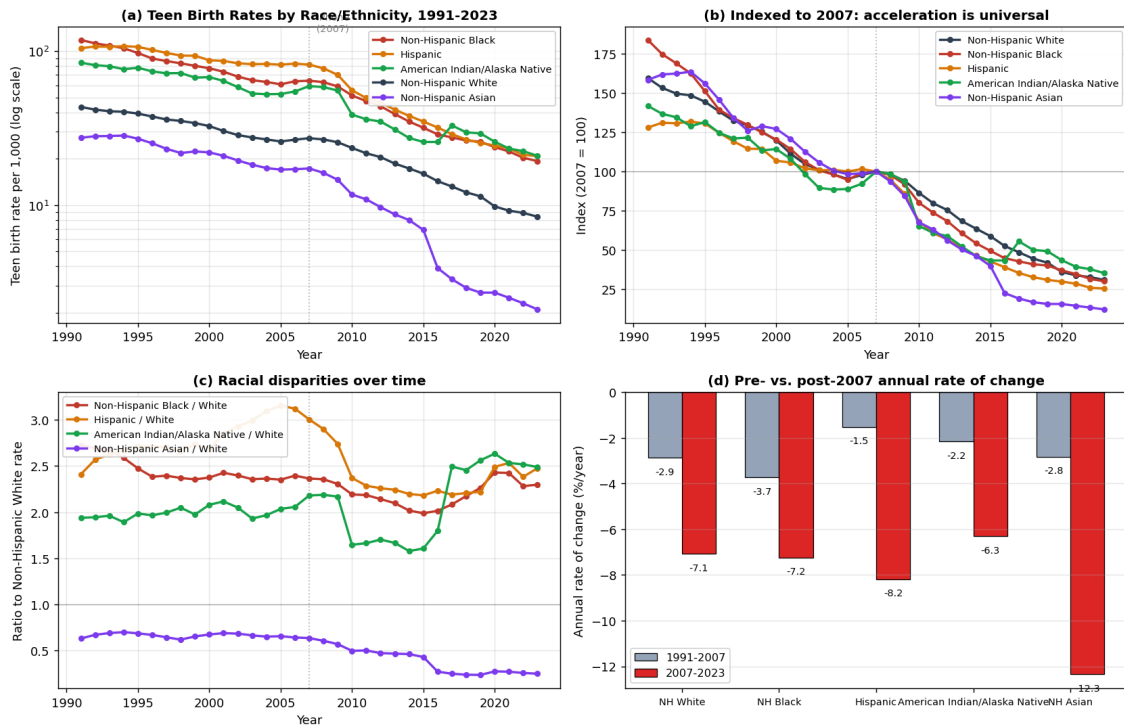


FIGURE 9. U.S. teen birth rates by race and Hispanic origin, 1991–2023. Panels: (a) levels on log scale; (b) indexed to 2007 = 100; (c) ratio to non-Hispanic White; (d) pre/post annualized rates of change. Source: NCHS NVSS.

densest.⁸

VI.I Dobbs as a Test of the Conception Channel

The *Dobbs* decision (June 2022) provides a secondary test of the conception-channel interpretation. The smartphone mechanism operates through conceptions, not through abortion access; *Dobbs* affects the resolution of conceptions but not the conception rate. If the conception channel is correct, the post-*Dobbs* reversal should be modest—bounded by the abortion ratio, which is 29–30 percent of teen pregnancies. The abortion-access story would predict a large reversal of the post-2007 decline.

⁸We measure the density tercile using the 2013 female 15–19 population to limit contamination from post-treatment migration; selective migration could still bias the gradient if movers differ systematically from stayers, though the direction of any such bias is not obvious. The point is descriptive: the gradient is what the mechanism would generate, and it would be hard to generate from contraceptive-access or welfare-reform stories that do not operate through peer-network density.

U.S. Teen Birth Rates by County Population Size, 2003-2020
3,135 counties, population-weighted tercile split by 2013 female 15-19 population

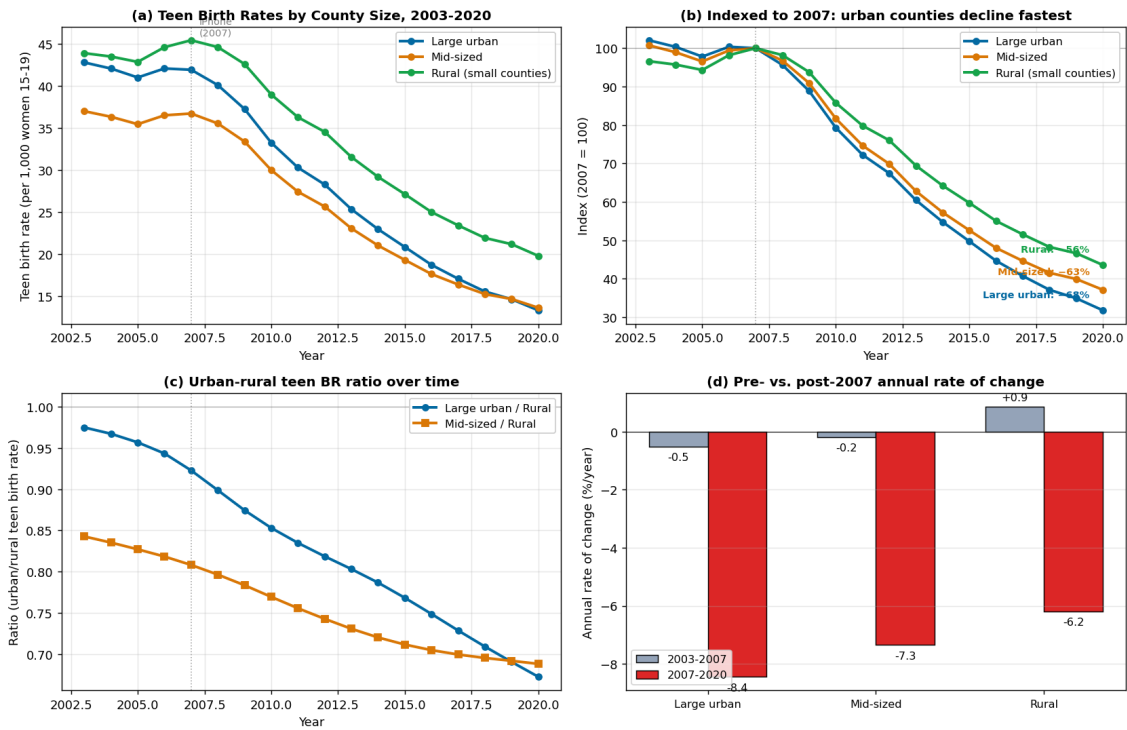


FIGURE 10. U.S. teen birth rates by county population size, 2003–2020. Counties sorted into three equal-population terciles. Source: NCHS Khan-Rossen county-level estimates.

Table 6 extrapolates the 2014–2019 log-linear trend in age-specific birth rates to 2024 and compares to observed 2024 rates.

TABLE 6. Post-Dobbs Birth Rate Deviations from Pre-Dobbs Trend Extrapolation

Age group	2014	2019	2024 actual	2024 trend	Gap
15–19	24.2	16.7	12.7	11.5	+10.2%
20–24	79.0	66.6	56.7	56.1	+1.0%
25–29	105.8	98.3	91.4	91.3	+0.1%
30–34	100.8	93.7	95.4	87.1	+9.5%
35–39	51.0	52.8	55.0	54.7	+0.6%
40–44	10.6	12.0	12.8	13.6	–5.8%

Notes: Rates per 1,000 women. 2024 trend is the 2014–2019 log-linear extrapolation. Source: NCHS NVSS.

The teen-15–19 reversal of +10.2 percent above pre-*Dobbs* trend is modest—about one-sixth of the total 2007–2019 smartphone-era decline of ~60 percent. Bounded by

the 30 percent abortion ratio, the conception channel predicts a teen reversal of at most 10–15 percent above trend; the observed 10.2 percent falls in this range. An abortion-driven account would predict a much larger reversal—50 percent or more of the post-2007 decline—which is not observed.

VII. INTERNATIONAL REPLICATION: ENGLAND AND WALES

The U.S. evidence in Section 6 is identified through within-country cross-sectional and within-county time variation, but the aggregate post-2007 acceleration is a common national shock whose magnitude cannot be separately identified from cross-sectional regressions alone. An international replication provides a second identification check: if the mechanism is a common global technology shock rather than a U.S.-specific policy confound, similar acceleration patterns should appear in other countries with comparable smartphone adoption trajectories but different healthcare and welfare environments. England and Wales provide a natural test case because (i) ONS administrative data measure conceptions directly rather than births, capturing the object $\lambda = \lambda_0 \cdot (1 - m)$ of Equation (4) without the need for an abortion-ratio placebo; (ii) Ofcom publishes local-authority-level 4G mobile coverage; and (iii) the UK’s universal National Health Service and Teenage Pregnancy Strategy environment is sufficiently different from U.S. Medicaid waivers, welfare reform, and the Affordable Care Act contraceptive mandate that any country-specific policy explanation cannot simultaneously fit both countries.

VII.A Data

Under-18 conception rates by local authority are from ONS Conception Statistics, Worksheet 6 (Office for National Statistics, 2023). Data cover 1998–2021 and report, for each lower-tier local authority in England, the number of conceptions to women aged under 18 and the rate per 1,000 women aged 15–17. A conception is defined as a pregnancy leading to either a live or still birth, or a legal abortion under the 1967 Abortion Act. The sample comprises four types of lower-tier local authority: Unitary Authorities (80), Non-metropolitan Districts (181), Metropolitan Districts (36), and London Boroughs (32).

After restricting to authorities with complete data for all 24 years, the balanced panel contains 294 local authorities \times 24 years = 7,056 observations.

Local-authority-level 4G mobile network coverage is from Ofcom’s Connected Nations reports for 2017–2020 (Ofcom, 2020). The key variable is the percentage of premises in each local authority with outdoor 4G signal from all four UK mobile network operators (EE, Three, O2, Vodafone). This all-operator coverage threshold is analogous to the FCC Form 477 broadband tier coding of Section 4, which measures the maximum tier of broadband available to a county’s population. Table 7 summarizes the distribution across years. The key feature for identification is the substantial 2017 cross-sectional variation (mean 87.4 percent, standard deviation 15.5 percentage points, range 24.2 to 100.0 percent), which compresses as the network approaches saturation by 2020.

TABLE 7. 4G Mobile Coverage Distribution by Year: % Premises with All-Operator Outdoor Coverage

	2017	2018	2019	2020
Local authorities	325	325	325	313
Mean	87.4	93.5	97.1	97.5
SD	15.5	8.4	4.4	3.8
Min	24.2	62.8	73.5	78.2
Max	100.0	100.0	100.0	100.0

Notes: Variable is all-operator outdoor 4G coverage from Ofcom Connected Nations mobile coverage data. LAs restricted to English E-code authorities matched to the ONS conception panel.

For the event-study specification, we use the standardized 2017 4G coverage level z_{ℓ}^{4G} as the cross-sectional treatment intensity measure, paralleling the U.S. paper’s use of a cross-sectional broadband index z_c . For the first-difference specification, we construct a continuous 4G coverage time series for each local authority: zero before 2013 (when UK 4G launched commercially), linearly interpolated between 2013 and the first observed value in 2017, observed values for 2017–2020, and the 2020 value carried forward to 2021.

Three features distinguish this analysis from the U.S. results of Section 6. First, the outcome variable is the under-18 conception rate rather than the 15–19 birth rate—a closer measure of the theoretical object. Second, the treatment variable is Ofcom 4G mobile coverage rather than FCC broadband tiers. Both are continuous and time-varying within

geographic units, but the English variable directly measures mobile network infrastructure rather than capturing it through a fixed-broadband proxy. Third, the panel is longer (24 years vs. 18) but narrower (294 local authorities vs. 3,136 counties).

VII.B National Trends

Figure 11 compares national teen fertility trends in England and the United States, indexed to 2007. Both countries exhibit a pre-2007 plateau followed by a sharp post-2007 acceleration, with near-identical timing. The English under-18 conception rate declined from 41.4 per 1,000 in 2007 to 13.0 in 2020, a fall of 69 percent. The U.S. teen birth rate declined from 42.5 to 15.3 over the same period, a fall of 64 percent.

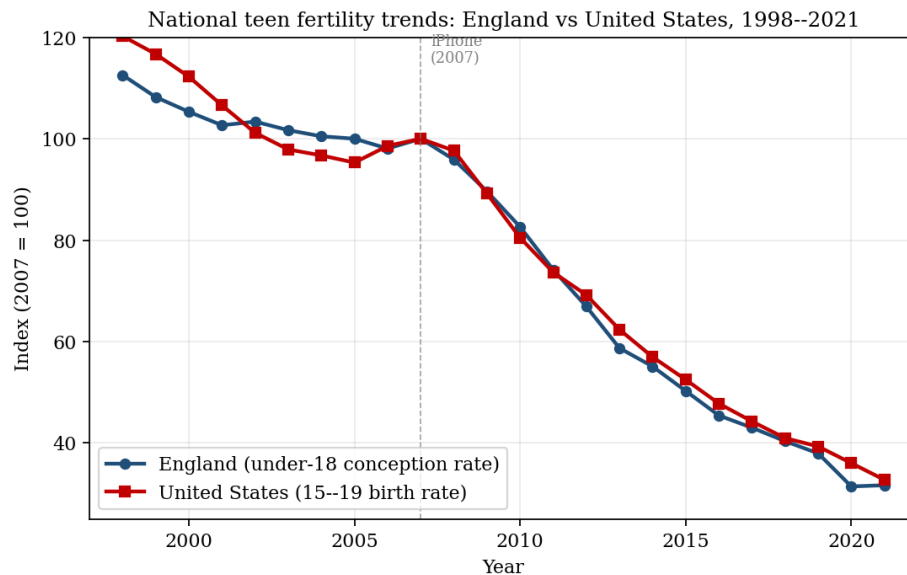


FIGURE 11. Teen fertility trends indexed to 2007, England vs. United States, 1998–2021. England series is ONS under-18 conception rate (per 1,000 women 15–17); U.S. series is NCHS age-specific birth rate (per 1,000 women 15–19). Source: ONS Conception Statistics, NCHS NVSS.

Table 8 formalizes the comparison. Pre-2007 annual log change was -1.3 percent per year in England versus -1.4 percent in the United States. Post-2007 annual log change accelerated to -8.9 percent per year in England and -7.9 percent in the United States. The acceleration ratio—post-2007 slope divided by pre-2007 slope—is $6.8\times$ for England and $5.6\times$ for the United States.

TABLE 8. National Teen Fertility Trends: England vs. United States

	England	United States
Measure	Conceptions (under 18)	Births (15–19)
Rate 2007	41.4	42.5
Rate 2020	13.0	15.3
Decline 2007–2020	–69%	–64%
Annual rate, 1998–2007	–1.3%/yr	–1.4%/yr
Annual rate, 2007–2020	–8.9%/yr	–7.9%/yr
Acceleration ratio	6.8×	5.6×

Notes: Annual rates are average annual log changes over the indicated period. England rates are ONS under-18 conceptions per 1,000 women aged 15–17. U.S. rates are NCHS births per 1,000 women aged 15–19. Source: ONS Conception Statistics; NCHS NVSS.

The near-identical timing across two countries with fundamentally different policy environments is difficult to reconcile with any country-specific policy explanation. UK teen-fertility policy over the post-2007 period operated through universal NHS contraceptive access and the Teenage Pregnancy Strategy; U.S. teen-fertility policy operated through Medicaid family-planning waivers, welfare reform, and after 2010 the Affordable Care Act contraceptive mandate. The near-simultaneous acceleration in both countries is consistent with a common global shock, of which the smartphone is the leading candidate.

VII.C Within-LA Distributed Lag

The UK panel merges ONS under-18 conception counts with Ofcom 4G outdoor coverage data for 294 local authorities over 1998–2021. As in the U.S. panel, we adopt a first-difference distributed-lag specification. The pre-trend diagnostic is weaker in England: the correlation between each local authority’s 1998–2007 log conception slope and its 2017 4G coverage is $r = -0.081$ ($p = 0.17$), against $r = -0.162$ in the U.S., so high-4G authorities in England were not on differential pre-2007 trajectories. The model still predicts a 3–5 year buildup that linear pre-trend corrections would absorb mechanically.

The FD specification is

$$\Delta \log(\text{CR}_{\ell,t}) = \gamma_t + \sum_{k=0}^3 \phi_k \Delta(4G_{\ell})_{t-k} + \varepsilon_{\ell,t}.$$

Table 9 reports the results. The contemporaneous effect is -0.00475 ($t = -2.56$): a one-percentage-point increase in 4G coverage is associated with a 0.48 percent decline in the teen conception rate in the same year. The third lag is also individually significant at -0.00374 ($t = -2.06$), consistent with the coordination model’s prediction of delayed equilibrium adjustment as peer networks accumulate. The cumulative four-year effect is -0.00566 , implying that a 10-percentage-point increase in 4G coverage reduces teen conceptions by approximately 5.7 percent over four years.

TABLE 9. Distributed-Lag First-Difference: Within-LA 4G Coverage Changes

	Coefficient	<i>t</i> -statistic
$\Delta(4G)_t$	-0.00475^{**}	-2.56
$\Delta(4G)_{t-1}$	$+0.00234$	$+1.31$
$\Delta(4G)_{t-2}$	$+0.00048$	$+0.34$
$\Delta(4G)_{t-3}$	-0.00374^*	-2.06
Cumulative 4-year	-0.00566	

Notes: 294 local authorities. Dependent variable is $\Delta \log(\text{CR}_{\ell,t})$. Treatment is the within-LA change in 4G all-operator outdoor coverage (percentage points). Cluster-robust SEs at the LA level. * $p < 0.05$, ** $p < 0.01$.

The lag structure differs from the U.S. paper’s monotonically negative pattern, with the first and second lags positive but insignificant. This may reflect measurement noise in the interpolated pre-2017 4G series, or differences in the timing of UK versus U.S. smartphone adoption dynamics. The cross-sectional correlation between 2017 4G coverage and the 2007–2020 log change in conception rates is $r = -0.195$ ($p < 0.001$): areas with higher 4G coverage experienced larger conception declines after 2007.

VII.D What the English Replication Adds

The English results parallel the U.S. results on every dimension that can be tested with available data, and add two pieces of substantive value. First, the conception measure records pregnancies regardless of outcome, embedding the abortion-access verification of Section 6.C by construction; the negative 4G coefficient on conceptions confirms the mechanism operates on the pregnancy-formation margin. Second, the UK’s NHS universal contraceptive access and absence of a welfare-reform cohort differ fundamentally from

the U.S. environment. That both countries exhibit near-identical post-2007 acceleration, and that within-country cross-sectional variation in mobile infrastructure independently explains the post-2007 decline in both settings, is difficult to reconcile with any single country-specific policy explanation.

VIII. BEYOND FERTILITY: A DIRECT TEST OF THE TIPPING PREDICTION

The coordination model makes a sharper empirical prediction than a smooth-transition story. In the coordination-tipping account two equilibria coexist until the price falls below the critical threshold \bar{p} , at which the in-person equilibrium ceases to exist and the peer network jumps to the digital equilibrium. Outcomes depending on in-person peer time should therefore exhibit synchronized *kinks*—changes in slope, not level—at the tipping date, with the sign of each kink determined by whether the outcome rises or falls with peer presence. Smooth dose-response cannot generate kinks; coordination tipping requires them. Fertility is one such outcome; this section tests the same prediction on five additional U.S. teen outcomes and on the cross-country teen suicide rate.

Figure 12 plots U.S. trajectories for the teen birth rate (NCHS), share of high-school students reporting having had sexual intercourse (YRBS), share of adolescents reporting a major depressive episode (NSDUH), teen suicide rate (CDC WONDER), and teen violent-crime arrest rate (FBI UCR), over a symmetric ten-year window around 2007. Each series kinks at 2007 in the predicted direction. The three outcomes that depend on in-person peer time all kink downward: birth-rate growth rate from -2.5 to -8.6 percent per year, sexual activity from -0.3 to -2.2 , violent-crime arrests from -3.5 to -6.2 . The two isolation-correlated outcomes reverse direction: adolescent depression from -3.9 to $+6.7$ percent per year, teen suicide from -2.4 to $+5.1$. Five outcomes, one date, all correctly signed—the empirical signature the coordination model predicts and the smooth-transition alternative does not.

The same kink prediction extends internationally. Figure 13 plots teen suicide rates from 1996 through 2018 for the 19 high-income countries used in Figure 1’s sub-panel A1,

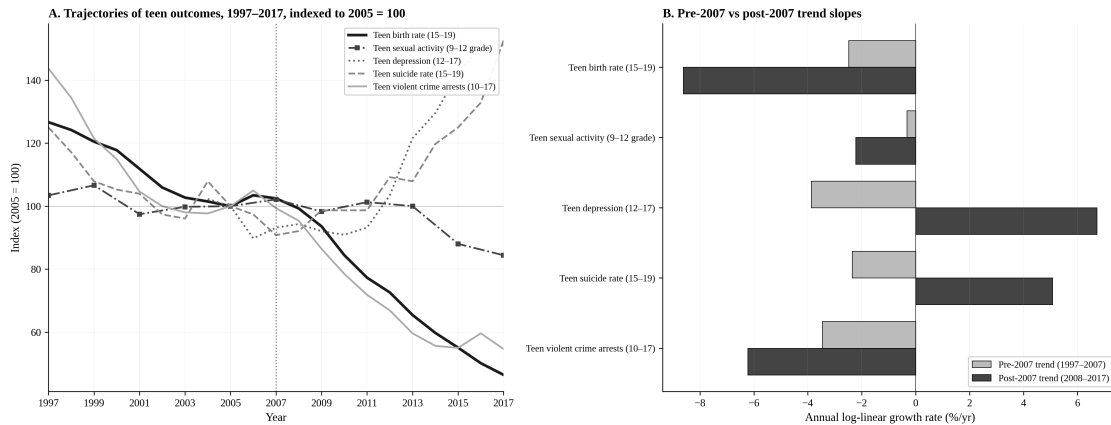


FIGURE 12. Synchronized 2007 kinks across five U.S. teen behavioral outcomes. **Panel A:** indexed trajectories, 1997–2017, with a vertical reference line at 2007. **Panel B:** pre-2007 (1997–2007) and post-2007 (2008–2017) annual log-linear growth rates. Three peer-time-dependent outcomes (birth rate, sexual activity, violent-crime arrests) kink to steeper declines; two isolation-correlated outcomes (depression, suicide) reverse from declining to rising. Sources: NCHS natality; YRBS biennial survey; NSDUH via CDC; CDC WONDER; FBI UCR via OJJDP Statistical Briefing Book.

indexed to each country’s 2007 value. In a smooth-transition account, the cross-country median trajectory should evolve gradually with no special feature at 2007. The data instead show a U-shape: the median declines steadily from roughly 135 in 1996 to its trough at 2007 and rises thereafter, ending at approximately 125 by 2018. The U-shape is the level-space manifestation of an upward kink in the suicide growth rate at 2007, mirroring the U.S. pattern in Figure 12. The United States shows the most pronounced U-shape; the United Kingdom a milder version of the same pattern; the interquartile range widens substantially after 2010, reflecting genuine cross-country heterogeneity in the post-2007 trajectory. Of the 20 high-income countries with sufficient data, 15 show post-2007 acceleration in teen suicide growth, with a median pre/post change of +4.4 percentage points per year. The international suicide pattern is therefore a partial replication of the U.S. kink: it confirms that the 2007 turning point is not specific to one country, but its magnitude varies more across countries than the fertility response in Figure 1.

The synchronized 2007 kink across U.S. teen outcomes, paralleled in the international teen suicide series, is what the coordination model predicts and what a smooth-transition story does not. This paper contributes the cross-country descriptive replication of the post-

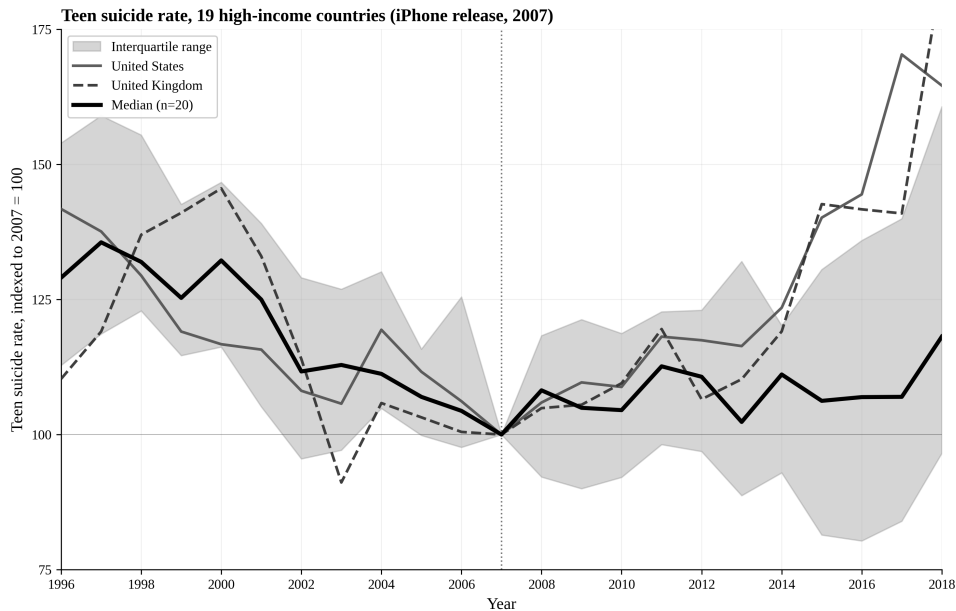


FIGURE 13. Teen suicide rate, 19 high-income countries, indexed to 2007 = 100. Cross-country median falls from approximately 135 in 1996 to 100 at 2007 and rises thereafter, ending at approximately 125 by 2018—a U-shape consistent with an upward kink in the suicide growth rate at the 2007 tipping date. Source: World Bank Data360 Human Capital Project, adolescent suicide rate per 100,000 population aged 15–19 (both sexes).

2007 break, the within-country fertility identification using terrain-ruggedness instruments, the coordination-equilibrium framework that nests the disparate findings as predictions of a single network-tipping shock, and the test of that framework’s distinctive prediction: a synchronized kink across multiple outcomes at a single tipping date. We extend that test to a direct IV exercise on teen suicide in the next subsection, and place the result alongside the existing causal mental-health literature in Section 8.2.

VIII.A A Direct IV for the Suicide Margin

The same identification strategy used in Table 3 for teen fertility can be applied directly to teen suicide. The empirical challenge is that CDC WONDER suppresses any county-cell with fewer than ten deaths, and direct queries of teen suicide deaths by county suppress the overwhelming majority of rural counties from the reportable sample.⁹

⁹We circumvent this limitation by recovering teen suicide counts via an all-age aggregate decomposition that yields a 1,899-county sample, an order of magnitude larger than the 363 counties available from a direct teen-suicide query. Partition external-cause deaths (ICD-10 V01–Y89) into a 2 × 2 grid by age (teen 15–19

We estimate a cross-sectional IV regression of the post-period teen suicide rate on average 2010–2019 broadband and 4G coverage, controlling for the pre-period suicide rate. Including the pre-period rate as a control on the right-hand side makes specifications with the post-period rate as outcome and with the pre–post change as outcome algebraically equivalent. The mechanism predicts a positive coefficient, opposite-signed to the negative fertility coefficient of Table 3: smartphone-driven substitution from in-person to digital peer time should reduce teen conceptions *and* raise teen suicide, both effects identified off the same instrument. Table 10 reports the progressive-controls build-up in the format of Table 3, with state fixed effects included throughout.

The IV estimates are positive and statistically significant in every column. The LTE coefficient ranges from +2.67 to +3.98 deaths per 100,000 per 10-percentage-point increase in mean 2010–2019 coverage (columns 4–6), and the broadband Tier 1 coefficient ranges from +3.21 to +6.04 (columns 1–3). Against a sample-mean post-period rate of 13.9 deaths per 100,000, a 10-percentage-point shift in coverage corresponds to a 19 to 44 percent increase in the teen suicide rate—of the same proportional order of magnitude as the LTE effect on teen fertility in Table 3, with first-stage *F*-statistics well above the weak-instrument threshold in every column.

The sign and magnitude of the IV estimates mirror the fertility result of Table 3 but in the opposite direction. Both effects are recovered from the same instrument applied to two behavioral outcomes that the coordination model predicts should move in opposite directions—fewer face-to-face encounters yield fewer conceptions, and more isolation from the in-person peer network yields more self-harm.

vs. all other ages) and cause (suicide X60–X84 vs. non-suicide external). Three queries per period—all-ages all-external, all-ages non-suicide-external, and not-15–19 suicide-only—fully identify all four cells, yielding teen suicide as $TS_c = (All \times All \text{ ext}) - (All \times Non\text{-sui ext}) - (Not\ 15\text{-}19 \times Suicide)$. Because each aggregate is much larger than teen suicide alone, each individually clears the ten-death threshold in nearly every U.S. county. We verify that population denominators reconcile across the queries and that recovered counts are non-negative. We further validate the decomposition against the prior teen-only subtraction (*teen all-external* minus *teen non-suicide external*); on the 1,207-county intersection where both methods deliver values, the implied counts agree exactly, county by county and period by period.

TABLE 10. Ruggedness IV for Digital-Infrastructure Effect on Teen Suicide Rate, Two Measures

	Dependent variable: Δ teen 15–19 suicide rate per 100,000, 1999–2007 to 2008–2020					
	Broadband penetration (per 10-pp)			LTE (4G) coverage (per 10-pp)		
	(1)	(2)	(3)	(4)	(5)	(6)
\hat{T}_c	+3.36** (1.38)	+3.21** (1.38)	+6.04*** (2.12)	+3.00** (1.24)	+2.67** (1.15)	+3.98*** (1.40)
First-stage F	82.1	72.8	38.8	53.3	55.4	43.0
First-stage coef. on $\log(\text{TRI} + 1)$	-0.144	-0.138	-0.097	-3.228	-3.303	-2.960
State fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Race shares	Yes	Yes	Yes	Yes	Yes	Yes
Bachelor's + share	Yes	Yes	Yes	Yes	Yes	Yes
Log total population	Yes	Yes	Yes	Yes	Yes	Yes
Log land area	Yes	Yes	Yes	Yes	Yes	Yes
Pre-period rate control	Yes	Yes	Yes	Yes	Yes	Yes
Religious composition (4 groups)	No	Yes	Yes	No	Yes	Yes
Republican vote share (2000–2020 mean)	No	No	Yes	No	No	Yes
Counties	1,899	1,899	1,899	1,899	1,899	1,899

Notes: 2SLS estimates. The estimated equation regresses the pooled post-period teen 15–19 suicide rate per 100,000 (2008–2020) on infrastructure coverage, controlling for the pre-period rate (1999–2007) on the right-hand side; the coefficient on \hat{T}_c is therefore algebraically identical to the coefficient one would obtain with the pre–post *change* in the suicide rate as the outcome, and is the quantity reported in the column header. Suicide counts are recovered through an all-age aggregate decomposition described in Section 8.1. Endogenous treatment T_c is the county 2010–2019 mean FCC Form 477 broadband Tier 1 share in columns (1)–(3) and the county 2010–2019 mean all-operator outdoor LTE (4G) coverage in percent of premises in columns (4)–(6). Coefficients are rescaled to a common *change in deaths per 100,000 per 10-percentage-point increase in the underlying infrastructure measure*: native broadband Tier 1 coefficient multiplied by 0.5; native per-pp LTE coefficient multiplied by 10. Standard errors in parentheses are robust to heteroskedasticity, also rescaled. The instrument is $\log(\text{TRI} + 1)$. Religion variables are adherence rates from the 2010 U.S. Religion Census (Grammich et al., 2012). Republican vote share is the county mean across 2000, 2004, 2008, 2012, 2016, and 2020 from MEDSL (MIT Election Data and Science Lab, 2024). Race shares are the non-Hispanic Black share and Hispanic share. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VIII.B The Mechanism’s Wider Footprint

The IV result on teen suicide complements the descriptive mental-health evidence in Figure 12 and the international suicide pattern in Figure 13, and sits alongside the existing causal evidence in the quasi-experimental literature. Braghieri, Levy, and Makarin (2022) use the staggered Facebook rollout across U.S. colleges and identify causal deteriorations in mental health; Donati et al. (2025) use 3G mobile broadband rollout in Italy and find increases in adolescent mental-health visits and antidepressant use. We treat the psychology literature (Twenge, Martin, and Campbell, 2018; Twenge et al., 2019; Haidt, 2024) as complementary on the breadth of outcomes. Causal identification of smartphone effects on depression, sexual activity, or violent-crime arrests requires either restricted-use mortality data with full county coverage or different identification strategies, and we leave these to future work.¹⁰ The descriptive kink-pattern in Figures 12 and 13 is nonetheless a model-driven test rather than a loose consistency check: kinks at the same date in the directions the model predicts are the discriminating signature of a coordination tipping point relative to a smooth-transition alternative.

IX. CONCLUSION

This paper argues that the post-2007 acceleration in young-adult fertility decline in the U.S. and U.K. is attributable to smartphones. The argument rests on five pieces of evidence: age-specific national trends showing a monotone gradient concentrated in the young-adult ages in which most unintended conceptions occur; cohort evidence showing the decline contains a quantum component that cannot be recovered later; a terrain-ruggedness IV identifying significant cross-sectional 2SLS effects of broadband and 4G, supplemented

¹⁰The county-year suppression rule that we circumvent for the cross-sectional change in suicide rates by pooling deaths over the full pre- and post-windows cannot be similarly relaxed for a within-county dynamic exercise. An annual county-year panel of teen suicide rates would require each county-year cell to clear the ten-death threshold, which retains only the largest urban counties and removes the rural variation that identifies the terrain-ruggedness instrument’s first stage. We therefore present only the cross-sectional change specification of Table 10 and do not estimate a corresponding panel event-study analogous to Section 6.6. Suicide classification is also subject to known endogenous reporting heterogeneity across coroner and medical-examiner jurisdictions (Rockett et al., 2014; Stone et al., 2018), which adds measurement noise that should attenuate the cross-sectional estimates we report.

by a within-county distributed-lag panel converging on a smaller dynamic effect with cointegration; time-use diary evidence showing teen in-person socializing fell from 68 to 38 minutes per day over 2003–2019 with year-by-year correlation $r = +0.97$ against the teen birth rate; and an English replication showing under-18 conception rates accelerated from -1.3 to -8.9 percent per year, with Ofcom 4G coverage yielding event-study and first-difference estimates that parallel the U.S. results. Additional confirmation of the mechanism comes from the suicide margin: applying the same terrain-ruggedness instrument to teen suicide rates recovers a positive effect of 4G coverage on the post-period suicide rate, opposite-signed to the fertility result and of similar proportional magnitude. Recovering symmetric, opposite-signed effects on two outcomes that the model predicts move in opposite directions in response to the same equilibrium shift is difficult to reconcile with any single-channel explanation.

The theoretical framework is a coordination model with positive network externalities in both in-person and digitally mediated interaction, driven by the falling hedonic price of the smartphone (down 84 percent in real terms 2007–2019). As the price falls through a critical threshold, the in-person equilibrium disappears and the economy tips. The shift is hysteretic: removing individual smartphones does not restore the in-person equilibrium because the peer network remains in the new one. The framing is that smartphones *accelerated* a decline already underway, rather than *caused* it ab initio; Kearney and Levine (2015) remain correct for the 1991–2005 period, and the paper’s contribution is to identify the force that explains why the decline did not stall at its 2005–2008 plateau.

REFERENCES

- Adams, Jessica A., Tamara S. Galloway, Debapriya Mondal, Sandro C. Esteves, and Fiona Mathews. “Effect of Mobile Telephones on Sperm Quality: A Systematic Review and Meta-Analysis.” *Environment International* 70 (2014): 106–112.
- Aizcorbe, Ana, David M. Byrne, and Daniel E. Sichel. “Getting Smart About Phones: New Price Indexes and the Allocation of Spending Between Devices and Services Plans in Personal Consumption Expenditures.” In *Measuring Economic Growth and Productivity*, edited by Barbara M. Fraumeni, 411–441. London: Academic Press, Elsevier, 2019.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. “The Skill Complementarity of Broadband Internet.” *Quarterly Journal of Economics* 130, no. 4 (2015): 1781–1824.
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. “The Welfare Effects of Social Media.” *American Economic Review* 110, no. 3 (2020): 629–676.
- Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton: Princeton University Press, 2009.
- Atasoy, Hilal. “The Effects of Broadband Internet Expansion on Labor Market Outcomes.” *Industrial and Labor Relations Review* 66, no. 2 (2013): 315–345.
- Bongaarts, John, and Griffith Feeney. “On the Quantum and Tempo of Fertility.” *Population and Development Review* 24, no. 2 (1998): 271–291.
- Braghieri, Luca, Ro’ee Levy, and Alexey Makarin. “Social Media and Mental Health.” *American Economic Review* 112, no. 11 (2022): 3660–3693.
- Branum, Amy M., and Jo Jones. “Trends in Long-acting Reversible Contraception Use Among U.S. Women Aged 15–44.” NCHS Data Brief No. 188, National Center for Health Statistics, February 2015.
- Chiu, Doris W., Isaac Maddow-Zimet, and Kathryn Kost. “Pregnancies, Births and Abortions in the United States, 1973–2020: National and State Trends by Age.” New York: Guttmacher Institute, 2024.
- Doepke, Matthias, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt. “The Economics of Fertility: A New Era.” In *Handbook of the Economics of the Family*, vol. 1, edited by Shelly Lundberg and Alessandra Voena. Amsterdam: Elsevier, 2023.
- Donati, Dante, Ruben Durante, Francesco Sobbrino, and Dijana Zejcirovic. “Lost in the Net? Broadband Internet and Youth Mental Health.” *Journal of Health Economics* 103 (2025): 103017.
- Grammich, Clifford, Kirk Hadaway, Richard Houseal, Dale E. Jones, Alexei Krindatch, Richie Stanley, and Richard H. Taylor. *2010 U.S. Religion Census: Religious Congregations & Membership Study*. Kansas City: Association of Statisticians of American Religious Bodies, 2012. County-level file distributed by the Association of Religion Data Archives, <https://www.thearda.com/data-archive?fid=RCMSCY10>.

- Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. “Engines of Liberation.” *Review of Economic Studies* 72, no. 1 (2005): 109–133.
- GSMA Intelligence. “Global LTE network forecasts and assumptions 2010–2015: 15 live LTE networks in December 2010 with 170 more planned by the end of 2015.” GSMA Intelligence research note, London, December 2010. Industry convention for feature-phone saturation threshold of 80 mobile cellular subscriptions per 100 inhabitants.
- Haidt, Jonathan. *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. New York: Penguin Press, 2024.
- Katz, Michael L., and Carl Shapiro. “Network Externalities, Competition, and Compatibility.” *American Economic Review* 75, no. 3 (1985): 424–440.
- Kearney, Melissa S., and Phillip B. Levine. “Investigating Recent Trends in the U.S. Teen Birth Rate.” *Journal of Health Economics* 41 (2015): 15–29.
- Kearney, Melissa S., Phillip B. Levine, and Luke Pardue. “The Puzzle of Falling US Birth Rates since the Great Recession.” *Journal of Economic Perspectives* 36, no. 1 (2022): 151–176.
- Kenny, Ryan P. W., Eugenie Evi Johnson, Adenike Adesanya, Catherine Richmond, Fiona Beyer, Carolina Calderon, Judith Rankin, Mireille B. Toledano, Maria Feychting, Mark S. Pearce, Dawn Craig, and Fiona Pearson. “The Effects of Radiofrequency Exposure on Male Fertility: A Systematic Review of Human Observational Studies with Dose–Response Meta-Analysis.” *Environment International* 190 (2024): 108817.
- Khan, Diba, Brady E. Hamilton, Lauren M. Rossen, Yulei He, Rong Wei, and Erin Dienes. “Teen Birth Rates for Age Group 15–19 in the United States by County, 2003–2020.” National Center for Health Statistics, 2024.
- Kim, Sungman, Donghyun Han, Jisun Ryu, Kichang Kim, and Yong Hyun Kim. “Effects of Mobile Phone Usage on Sperm Quality—No Time-Dependent Relationship on Usage: A Systematic Review and Updated Meta-Analysis.” *Environmental Research* 202 (2021): 111784.
- Kolko, Jed. “Broadband and Local Growth.” *Journal of Urban Economics* 71, no. 1 (2012): 100–113.
- Manuelli, Rodolfo E., and Ananth Seshadri. “Frictionless Technology Diffusion: The Case of Tractors.” *American Economic Review* 104, no. 4 (2014): 1368–1391.
- Milgrom, Paul, and John Roberts. “Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities.” *Econometrica* 58, no. 6 (1990): 1255–1277.
- MIT Election Data and Science Lab. “County Presidential Election Returns 2000–2020.” Harvard Dataverse, 2024. <https://doi.org/10.7910/DVN/V0QCHQ>.
- Ofcom. “Connected Nations 2020.” London: Office of Communications, 2020. Available at <https://www.ofcom.org.uk/siteassets/resources/documents/research-and-data/multi-sector/infrastructure-research/connected-nations-2020/connected-nations-2020.pdf>.

- Office for National Statistics. “Conception Statistics, England and Wales: Reference Tables, 2021.” Statistical Bulletin, Office for National Statistics, 2023. Available at <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/conceptionandfertilityrates/datasets/conceptionstatisticsenglandandwalesreferencetables>.
- Rahban, Rita, Alfred Senn, Serge Nef, and Martin Röösl. “Association between Self-Reported Mobile Phone Use and the Semen Quality of Young Men.” *Fertility and Sterility* 120, no. 6 (2023): 1181–1192.
- Raine, Tina R., Cynthia C. Harper, Corinne H. Rocca, Richard Fischer, Nancy Padian, Jeffrey D. Klausner, and Philip D. Darney. “Direct Access to Emergency Contraception Through Pharmacies and Effect on Unintended Pregnancy and STIs: A Randomized Controlled Trial.” *JAMA* 293, no. 1 (2005): 54–62.
- Raymond, Elizabeth G., James Trussell, and Chelsea B. Polis. “Population Effect of Increased Access to Emergency Contraceptive Pills: A Systematic Review.” *Obstetrics & Gynecology* 109, no. 1 (2007): 181–188.
- Riley, Shawn J., Stephen D. DeGloria, and Robert Elliott. “A Terrain Ruggedness Index That Quantifies Topographic Heterogeneity.” *Intermountain Journal of Sciences* 5, no. 1–4 (1999): 23–27.
- Rockett, Ian R. H., Eric D. Caine, Hilary S. Connery, and Haomiao Jia. “Mortality and Risk Factors Associated with the Misclassification of Suicide as an Undetermined or Unintentional Injury Death.” *Suicide and Life-Threatening Behavior* 44, no. 6 (2014): 718–728.
- Romero, Lisa, Karen Pazol, Lee Warner, Shanna Cox, Charlan Kroelinger, Ghenet Besera, Amanda Brittain, Tanya Telfair Lebherz, and Wanda Barfield. “Reduced Disparities in Birth Rates Among Teens Aged 15–19 Years—United States, 2006–2007 and 2013–2014.” *MMWR Morbidity and Mortality Weekly Report* 65, no. 16 (2016): 409–414.
- Schelling, Thomas C. *Micromotives and Macrobehavior*. New York: W. W. Norton, 1978.
- Stone, Deborah M., Thomas R. Simon, Katherine A. Fowler, Scott R. Kegler, Keming Yuan, Kameron J. Holland, Asha Z. Ivey-Stephenson, and Alex E. Crosby. “Vital Signs: Trends in State Suicide Rates—United States, 1999–2016 and Circumstances Contributing to Suicide—27 States, 2015.” *MMWR Morbidity and Mortality Weekly Report* 67, no. 22 (2018): 617–624.
- Twenge, Jean M., Gabrielle N. Martin, and W. Keith Campbell. “Decreases in Psychological Well-Being Among American Adolescents After 2012 and Links to Screen Time During the Rise of Smartphone Technology.” *Emotion* 18, no. 6 (2018): 765–780.
- Twenge, Jean M., A. Bell Cooper, Thomas E. Joiner, Mary E. Duffy, and Sarah G. Binau. “Age, Period, and Cohort Trends in Mood Disorder Indicators and Suicide-Related Outcomes in a Nationally Representative Dataset, 2005–2017.” *Journal of Abnormal Psychology* 128, no. 3 (2019): 185–199.

U.S. Department of Agriculture, Economic Research Service. “Area and Road Ruggedness Scales.” Washington, DC: USDA ERS, 2024. Available at <https://www.ers.usda.gov/data-products/area-and-road-ruggedness-scales>.

Yu, Gang, Zhiming Bai, Chao Song, Qing Cheng, Gang Wang, Zeping Tang, and Sixing Yang. “Current Progress on the Effect of Mobile Phone Radiation on Sperm Quality: An Updated Systematic Review and Meta-Analysis of Human and Animal Studies.” *Environmental Pollution* 282 (2021): 116952.