

WIDE AND SHALLOW: DIGITAL TECHNOLOGY AND THE POST-2007 FERTILITY DECLINE

NATHAN HUDSON AND HERNAN MOSCOSO BOEDO[†]

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

May 14, 2026

ABSTRACT

Fertility has fallen sharply across the high-income world since 2007. For the first time in human history, world fertility is close to the replacement rate. We propose a new mechanism that accounts for this decline: the digital revolution has fundamentally reshaped how humans interact with one another, favoring broad and shallow connections at the expense of the deeper ones that require sustained in-person investment. We formalize this through a dynamic household-production model in which relational capital affects partnership formation, divorce, and fertility. Calibrated to US 2007–2024 time-use and fertility data, the model accounts for the observed fertility decline; a counterfactual holding the relative price of phones at its 2007 level explains 43 percent of the change in fertility rates in the US. The same calibration matches the median OECD country’s observed decline in total fertility rate. Two causal designs identify the mechanism: a terrain-ruggedness IV on county-level broadband and 4G LTE rollout for teen fertility, and a cohort-Bartik design on 4G LTE rollout for partnership formation among smartphone-era cohorts.

Keywords: fertility, partnership formation, digital technology, leisure, social capital, household production, second demographic transition.

JEL Codes: J11, J12, J13, O33, D13, Z13.

I. INTRODUCTION

Fertility has collapsed below replacement nearly everywhere in the high-income world, with a sharp acceleration beginning around 2007. For the first time in human history, world

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

fertility is close to the replacement rate. The OECD median total fertility rate fell from 1.71 in 2007 to 1.44 in 2023 (Figure 1); pre-2007 fertility had been roughly stable for two decades. The acceleration appears across countries that differ widely in income, housing costs, childcare access, female labor force participation, family policy, and religion. The literature has documented this pattern carefully and identified a range of contributing factors (Kearney, Levine, and Pardue, 2022; Doepke et al., 2023); we propose a new mechanism that accounts for most of this decline.




FIGURE 1. Total fertility rate, 1990–2023, all 38 OECD member countries. Light gray lines: each country’s individual TFR trajectory. Gold shaded band: inter-quintile range across the 38 countries, computed year-by-year. Gold-diamond line: cross-country median. Navy line: United States. Vertical dotted line marks the June 2007 launch of the iPhone. Horizontal dotted line marks the replacement rate of 2.1. The post-2007 acceleration is visible in the median (which had been stable at ~ 1.7 and falls to 1.44 by 2023), in the contraction of the inter-quintile band, and in the US trajectory (at replacement through 2007, then steep decline). *Source:* World Bank World Development Indicators, series SP.DYN.TFRT.IN (World Bank, 2024).

The mechanism we propose is that the digital revolution has fundamentally reshaped how humans interact with one another, favoring broad and shallow connections at the expense of the deeper ones that require sustained in-person investment. Households allocate time between two kinds of social activity: broad, shallow connections that digital

technology makes nearly free, and deep, durable connections that require sustained in-person investment. The second kind of connection is the input to partnership formation, marital stability, and chosen fertility. As digital technology reallocates household time from the second to the first, deep relationships erode, partnerships form less often, the partnerships that do form are weaker, and conditional fertility falls. The same mechanism delivers the joint pattern observed in the data — collapsing in-person socializing, rising digital leisure, falling marriage rates, age-differential divorce dynamics, and the post-2007 fertility break.

We formalize this through a dynamic household-production model in which a single depth stock enters partnership formation, divorce, and fertility jointly. The model is calibrated to US time-use and fertility data over 2007–2024 with five parameters identified by six time-series moments. The fit accounts for essentially all of the observed kids-per-adult decline. A counterfactual that holds the relative price of phones at its 2007 level produces 43 percent of the actual decline — large enough that the digital channel is a first-order driver of the post-2007 break, but small enough that other channels carry the remainder. The same calibration applied country-by-country, with country-specific wage and price paths, matches the median OECD country’s observed total fertility rate decline to within one percentage point.

The mechanism is identified causally within this paper at two age windows. At the youngest end (Section VI), a terrain-ruggedness IV on county-level broadband and 4G LTE rollout shows that 10 additional percentage points of long-run digital infrastructure caused a 9.8 percentage point larger decline in 15–19 birth rates over 2003–2018 (first-stage $F = 82$), with a parallel 18.9 percentage point effect for broadband ($F = 73$). At the partnership-formation margin (Section VII), a cohort-Bartik design that interacts terrain ruggedness with cohort-level 4G LTE exposure during the 18–25 window yields a 0.34 reduction in the cohort female ever-married rate at full LTE saturation across the 1985–1995 smartphone-era cohorts (3,129 US counties; first-stage $F = 122$); the effect is robust to opioid, geographic, and cohort-interacted pre-period-education controls. Just as the first demographic transition was a quantity-to-quality reallocation on children, induced

by industrialization that lowered the relative price of child quality (Galor and Weil, 2000; Doepke et al., 2023), the second demographic transition is a quality-to-quantity reallocation on connections, induced by digitalization that lowered the relative price of breadth.

I.A The evidence we are after

Seven empirical regularities motivate the model and constitute the targets of our calibration. In order: (i) a global, simultaneous post-2007 acceleration in the fertility decline; (ii) reallocation of time toward digital leisure and away from in-person socializing; (iii) declining partnership formation and rising solo-living rates among young adults; (iv) deterioration in the depth of partnerships that do form, visible in confidant counts, loneliness indices, and self-reported relationship-quality measures; (v) a digital substitution path for couple formation that displaced friends and family networks; (vi) compositional selection in divorce dynamics; and (vii) a cross-country mapping from digital-saturation timing to fertility-decline timing. The first six constitute the targets of our calibration; the seventh provides the cross-country identifying variation we exploit in Section V.G.

Fact 1: Global, simultaneous post-2007 acceleration. Period total fertility rates show a discrete break around 2007–2010 in nearly every high-income country, not a continuation of pre-existing trends. Pre-2007 fertility had been roughly stable across the panel for two decades. Post-2010 fertility falls monotonically in almost every country, with the steepest declines concentrated where digital saturation rose fastest (Kearney and Levine, 2025); convergence below replacement is occurring across countries that differ in essentially every conventional explanatory dimension (Kearney, Levine, and Pardue, 2022). Table 1 summarizes the cross-OECD distribution: the median country’s TFR fell from 1.71 in 2007 to 1.44 in 2023; the median country-level decline (across the 38 OECD members) was –18.7 percent, and the count below replacement rose from 32 of 38 OECD members to 37 of 38.

Panel A of Figure 2 shows the full panel of 217 countries from 1970 onwards. The first demographic transition — the long fall in fertility toward replacement — is visible across every income group. World total fertility stood at 2.20 in 2023 and on current trends

TABLE 1. Total fertility rate across the 38 OECD member countries, 2007 vs. 2023. Median country crossed from 1.71 to 1.44 over the window; 37 of 38 OECD members were below the 2.1 replacement rate by 2023, up from 32 of 38 in 2007. Range across the panel in 2023 runs from South Korea at 0.72 to Israel at 2.85.

	TFR 2007	TFR 2023	Median % change [†]
OECD median (across 38)	1.71	1.44	−18.7%
OECD 20th percentile	1.37	1.25	−25.7%
OECD 80th percentile	2.00	1.55	−6.0%
United States	2.12	1.62	−23.8%
OECD weighted average (population-weighted)	1.73	1.49	−13.9%

Sources: World Bank World Development Indicators (SP.DYN.TFRT.IN) for country-level TFR; OECD Family Database SF2.1 (OECD, 2024) for OECD weighted average. Percentile rows in columns 1–2 are computed across the unweighted distribution of 38 OECD member-country TFR values year-by-year. [†]The percent-change column reports the relevant percentile of the distribution of country-level percent changes between 2007 and 2023, which is a different statistic from the change implied by the percentile levels in columns 1–2.

will cross below replacement within the present decade.

The level plot understates how sharp the recent break has been; Panel B of Figure 2 shows the same panel transformed into annualized growth rates. High-income countries experienced a fertility *recovery* between roughly 2002 and 2010 — the only sustained period of positive fertility growth in the post-1970 high-income panel — that reversed sharply after 2010, with the high-income growth rate falling to approximately −2 percent per year by 2020. The pre-2007 trajectory was not a deepening continuation of the demographic transition; it was a tentative recovery that reversed at exactly the timing of mass smartphone adoption.

Fact 2: Time use shifted from in-person socializing to digital leisure, with the family-formation cohorts hit hardest. We compute time-use trends directly from the American Time Use Survey activity-summary microdata covering 2003 to 2024 (n=252,808 respondent-days), applying ATUS final weights. Among all adults aged 15+, in-person socializing (ATUS code 120101) fell from 40.6 to 31.8 minutes per day (−22 percent), and games plus computer-based leisure (codes 120307, 120308) rose from 17.6 to 34.0 minutes per day (+93 percent). Figure 3 stratifies the series by age. Among the 15–24

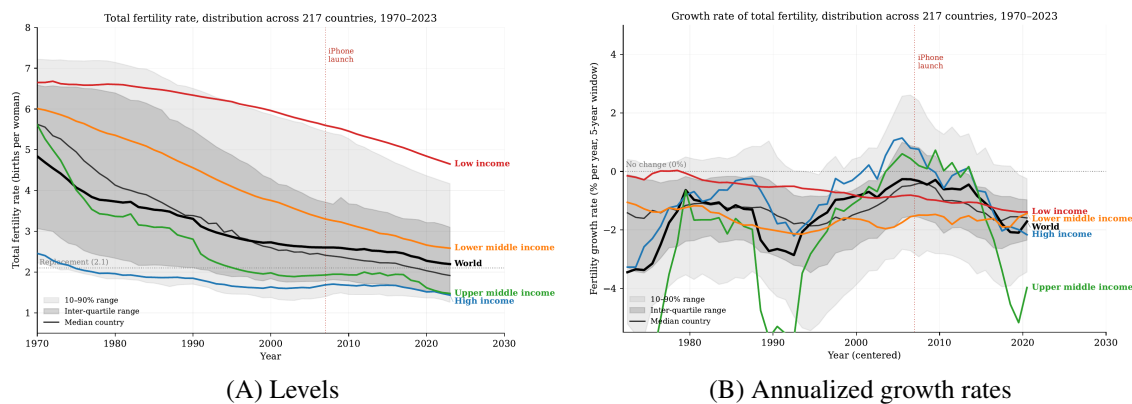


FIGURE 2. Total fertility rate, 1970–2023, distribution across 217 countries. Panel A: levels. Light shaded area: 10th–90th percentile range across countries; darker shaded area: inter-quartile range; thin dark line: cross-country median. Horizontal dotted line: replacement rate of 2.1. Panel B: centered five-year annualized log change in TFR, percent per year. In both panels, colored lines show World Bank income-group aggregates (population-weighted) and the vertical dotted line marks the 2007 iPhone launch. *Source:* World Bank World Development Indicators, series SP.DYN.TFRT.IN; authors’ calculations for Panel B.

cohort, in-person socializing fell 34 percent and digital leisure rose 86 percent; among the 25–34 cohort — the prime family-formation years — digital leisure tripled (16 to 50 minutes per day). Time spent on recreational computing and gaming rose by an average of 99 hours per year for young men between 2004–2007 and 2012–2015, accounting for roughly three-quarters of their total leisure increase (Aguiar et al., 2021). Comparable shifts appear in European Harmonised Time Use Survey data (Eurostat, various); for adolescents specifically, in-person socializing roughly halved while digital leisure roughly tripled (Section VI.A).

Fact 3: Partnership formation collapsed on the extensive margin (US). Marriage rates and partnership rates fell sharply across the same post-2007 window. In the US, the share of adults aged 40 who had never been married rose from 6 percent in 1980 to 25 percent in 2021 (Pew Research Center, 2023); among current US adults, 86 percent of those aged 18–24 and 42 percent of those aged 25–39 are unpartnered (Pew Research Center, 2025). The median age at first marriage rose from 26.8 (men) and 25.1 (women) in 2000 to roughly 30 and 28 in 2023 (U.S. Census Bureau, 2023). Comparable declines appear across the high-income panel, though the extensive-versus-intensive split varies:

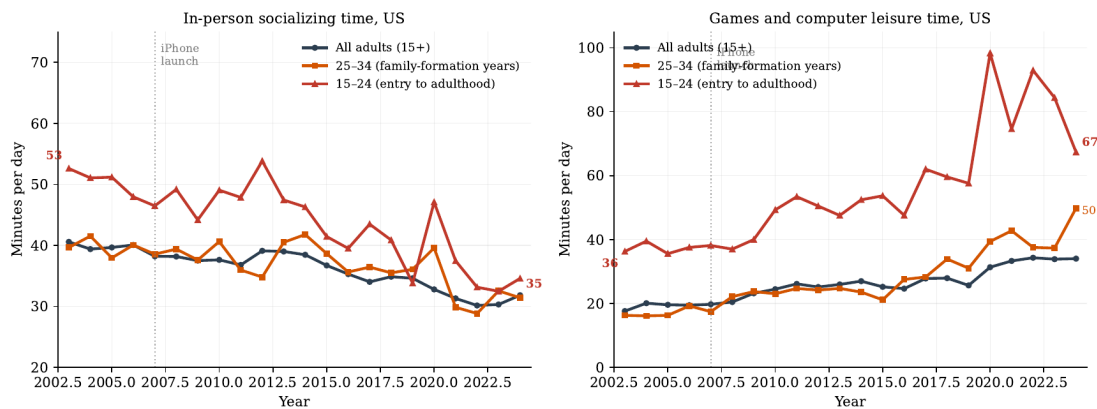


FIGURE 3. US time use by age cohort, 2003–2024, computed directly from the American Time Use Survey activity-summary microdata (atussum 0324, BLS) using ATUS final weights TUFNWGTP. Left panel: average minutes per day in in-person socializing (ATUS code 120101). Right panel: average minutes per day in games and computer-based leisure (ATUS codes 120307 and 120308). Three series shown: all adults aged 15+, the 25–34 family-formation cohort, and the 15–24 cohort entering adulthood. The vertical dotted line marks the 2007 iPhone launch. Among the 15–24 cohort, in-person socializing fell from 53 to 35 minutes per day (–34 percent) while games and computer leisure rose from 36 to 67 minutes per day (+86 percent). Among the 25–34 cohort, games and computer leisure rose from 16 to 50 minutes per day — a tripling over the period. *Source:* authors’ weighted calculations from BLS American Time Use Survey activity-summary file, 2003–2024 pooled (n=252,808 respondent-days).

Hellstrand, Nisén, and Myrskylä (2022) find roughly 75 percent of Finland’s post-2010 first-birth decline is from within-union fertility rather than union formation, the opposite weighting from the US, while Yoo (2023) finds the Korean decline shifts across cohorts from intensive to extensive to within-marriage childlessness. Partnership decline precedes fertility decline within cohorts (Kearney and Levine, 2025). Figure 4 stratifies US adults 25–44 by relationship status: the share married fell from 64 to 54 percent, the no-partner share rose from 29 to 35 percent, and conditional on marriage the mean number of household children fell from 1.66 to 1.55. A simple decomposition attributes 70 percent of the kids-per-adult decline to the share-married (extensive) margin and 30 percent to the children-conditional-on-marriage (intensive) margin; the structural model in Section III produces decline through a depth stock that supports both margins simultaneously.

Fact 4: Where partnerships form, they are weaker. General Social Survey data show the average number of confidants with whom Americans discuss important matters

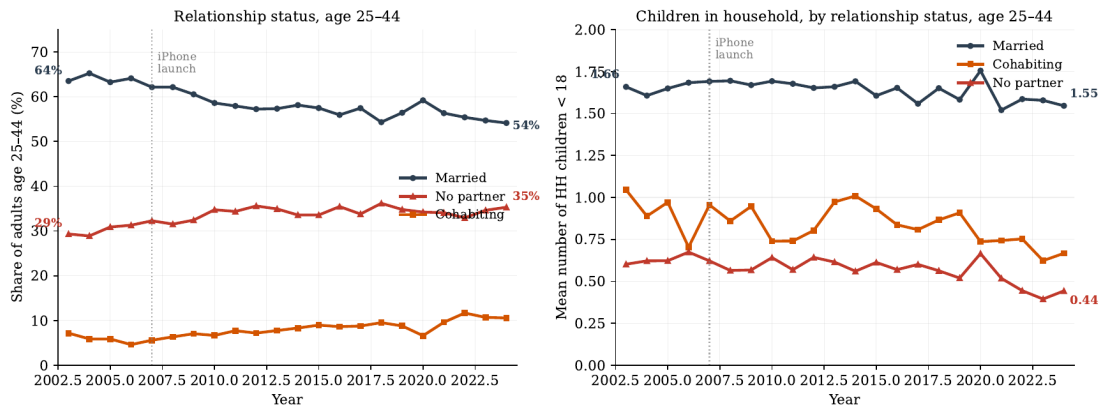


FIGURE 4. Relationship status and household children, US adults age 25–44, 2003–2024, weighted ATUS microdata. Left panel: share of the prime-childbearing-age population in each relationship category. Right panel: mean number of household children under 18 per adult, conditional on relationship status. The vertical dotted line marks the 2007 iPhone launch.

fell from approximately three in 1985 to two in 2004, with the share reporting zero confidants more than doubling (McPherson, Smith-Lovin, and Brashears, 2006). The Cigna Loneliness Index has reported persistent rises in reported loneliness across the post-2010 period despite expanding nominal connectedness through social platforms (Cigna, 2020). Meanwhile the number of digital “connections” — followers, friends, contacts — has grown by orders of magnitude. Rising breadth, falling depth: the empirical signature of the connection quantity-quality reallocation we model.

Fact 5: Divorce dynamics reveal compositional selection on depth. At first glance the US divorce data appear to contradict our depth-erosion claim: the refined divorce rate fell from its 1980 peak of 22.6 per 1,000 married women to 14.2 in 2024 (Loo, 2024; Allred, 2025a). Closer inspection reveals compositional selection on the marriage entry margin. Among women under 45 — the cohorts most affected by the post-2007 partnership-formation collapse — divorce rates have fallen substantially (33.0 to 21.4 per 1,000 for the 15–24 cohort, 31.4 to 22.0 for 25–34, 22.4 to 18.0 for 35–44 (Allred, 2025b)). Among older cohorts they have moved in the opposite direction: rates more than doubled for women 55–64 and rose nearly fivefold for women 65+; the share of all divorces involving someone aged 50+ rose from 8.7 percent in 1990 to 36 percent by 2019 (Brown and Lin, 2022). Figure 5 shows both patterns.

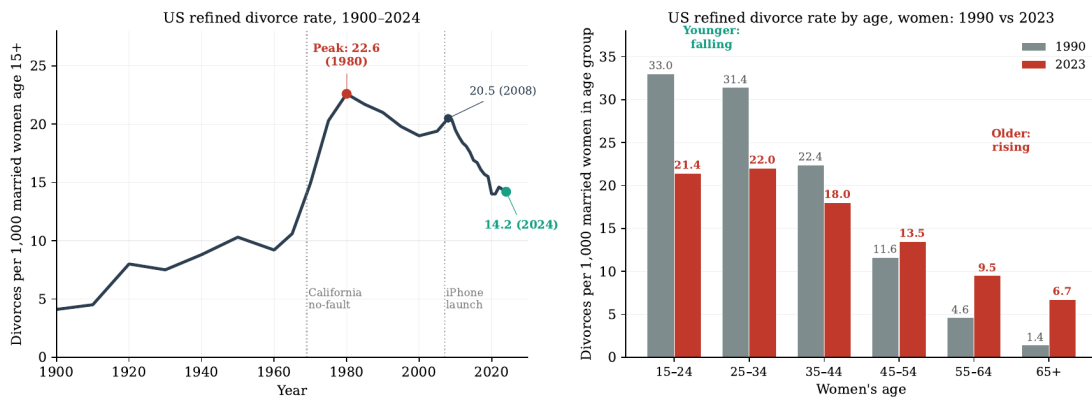


FIGURE 5. US divorce rate evolution. Left panel: refined divorce rate, 1900–2024. Right panel: refined divorce rate by women’s age group, 1990 and 2023. *Sources:* NCFMR Family Profiles FP-24-11, FP-25-31, FP-25-24, based on NCHS National Vital Statistics 1900–2000 and US Census Bureau American Community Survey 2008–2024.

The marriage-formation side mirrors this pattern. First marriage rates among the 18–29 cohort fell roughly 45 percent for both women and men between 1990 and 2019 (Brown, Lin, and Mellencamp, 2023), while at older ages they rose substantially (more than 125 percent for both sexes aged 50–59); the aggregate refined marriage rate fell 59 percent over 1970–2024 (Westrick-Payne, 2025). Figure 6 reports the by-age detail.

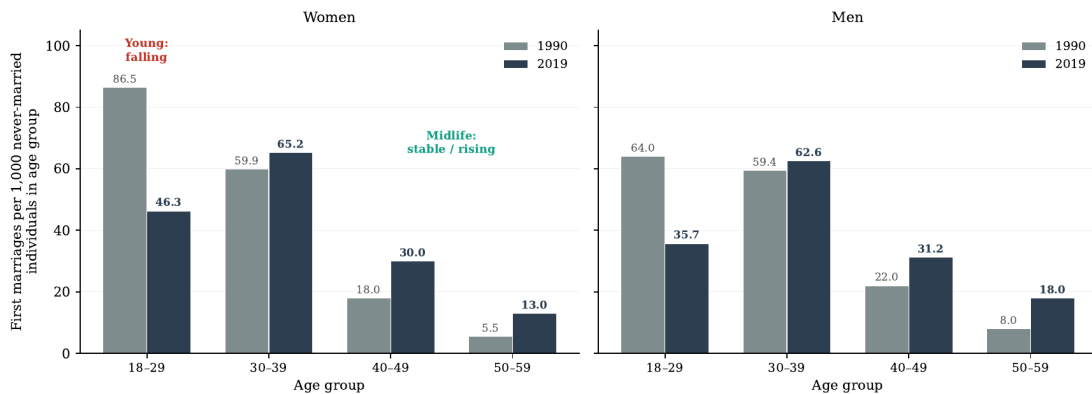


FIGURE 6. US first marriage rate by age group and sex, 1990 and 2019, per 1,000 never-married individuals. The pattern mirrors the age-stratified divorce-rate pattern in Figure 5. *Source:* Brown, Lin, and Mellencamp (2023). Aggregate rate: Westrick-Payne (2025).


Considered jointly, the divorce-rate decline among adults under 45 is the survival selection induced by the partnership-formation collapse: as marriage contracts on the extensive margin, the marriages that do form are increasingly composed of partnerships

with above-threshold depth stocks. The simultaneous rise in gray divorce among adults over 50 is the depth-erosion margin operating on the long-marriage stock. A simple “digital technology weakens relationships” story would predict rising divorce at all ages; a simple “matching efficiency” story would predict rising divorce only among the recently married. Neither matches; the selection-and-erosion structure of our framework delivers both panels of Figure 5 jointly.

Fact 6: Causal validation on the matching channel, and the rise of online relationship origination. Recent quasi-experimental evidence shows that digital matching technology does not produce more partnerships. Büyükeren et al. (2026) use Tinder’s college rollout via Greek organizations and find a sharp rise in casual sexual activity but no effect on long-term partnership formation. Section VI.A below provides parallel causal evidence on the teen margin using a terrain-ruggedness instrument. Meanwhile, the matching infrastructure has shifted dramatically online: among US heterosexual couples first meeting in 1994–1998, 3.9 percent met online; by 2017 the share had reached 39 percent, surpassing meeting through friends around 2013 (Rosenfeld, Thomas, and Hausen, 2019; Rosenfeld and Hausen, 2024). The Knot (2024) report 27 percent of US 2024 newlyweds met on a dating app specifically, with broader online channels accounting for an estimated 60 percent. Figure 7 shows the long historical sweep.

Pathways that historically required social-network embeddedness — friends and family — have lost half their mid-1990s share, while online matching has displaced both as the modal channel. The post-2009 acceleration of online origination maps onto the same digital-saturation window driving depth-stock erosion, and *coincides* with declining partnership formation rather than offsetting it.

Fact 7: Cross-country digital saturation timing maps onto fertility decline timing. The first iPhone launched in June 2007. Smartphone penetration crossed 50 percent of the population at different dates across the panel: roughly 2012 in the United States and South Korea, 2013–2014 in Japan and Western Europe, and later in Southern Europe (Newzoo, 2021; Statista, 2024). By 2023, smartphone penetration had reached 86 percent in North America, 82 percent in Western Europe, and 98 percent of South Korean adults



Source: Rosenfeld, Thomas, and Hausen (2019), PNAS 116(36): 17753–17758, Figure 1. LOWESS-smoothed shares from the HCMST 2017 nationally representative survey of US heterosexual couples (n=2,462). Categories sum to more than 100% because friends, family, and coworkers can overlap with other categories.

FIGURE 7. Share of US heterosexual couples who first met their partner via each channel, by year of meeting, 1940–2017. *Source:* Authors’ reconstruction from Rosenfeld, Thomas, and Hausen (2019) Figure 1, Stanford HCMST 2017 (n=2,462), LOWESS-smoothed.

(Statista, 2024). The cross-country timing of this saturation correlates strongly with the cross-country timing of the post-2007 fertility break: countries with earlier and faster smartphone saturation experienced earlier and steeper fertility declines, conditional on standard controls. This cross-country timing variation, together with the within-country cohort variation in age at smartphone saturation, provides our main source of identifying variation in Section IV.

These seven facts cannot be jointly explained by any standard fertility framework. They can be jointly explained by a single mechanism: digital technology shifting the relative price and relative utility of broad-shallow versus deep connection production, with downstream consequences for partnership formation and fertility through household production.

I.B The main forces we capture

The model formalizes four forces. *Time reallocation*: digital leisure has higher utility per minute than the time investments that build relationship depth, producing a reallocation away from depth-producing activities even without preference shifts (Aguiar et al., 2021; Kearney and Levine, 2025). *Asymmetric digital substitution*: the digital good substitutes for some relationship functions (identity validation, ambient companionship, stimulation) but not others (practical cooperation, embodied attachment, risk-sharing, the reproductive option), entering the time-share input but not the depth stock.² *A depth stock that depreciates*: relationship depth is a stock produced by ongoing time investment, the binding input to both partnership formation and conditional fertility, eroding endogenously when investment reallocates away. *A discrete partnership entry decision*: conditional on depth, agents make a discrete entry choice with a threshold determined by alternatives. The discrete margin matters in principle, but at the calibrated heterogeneity it does not bind: the depth-stock channel alone delivers the aggregate magnitude (Section V.B). These four forces, with heterogeneous agents, reproduce the cross-country and cross-cohort fertility patterns observed since 2007.

I.C Roadmap

Section II situates the paper relative to the existing literatures on fertility, leisure technology, marriage market frictions, and social capital. Section III presents the model. Section IV describes the calibration. Section V presents the quantitative results and counterfactuals, including the cross-country reproduction. Section VI adds within-country reduced-form evidence on age-specific fertility. Section VII provides causal IV identification on the partnership-formation margin. Section VIII concludes.

²The asymmetry organizes the connection-quality literature (Cacioppo and Patrick, 2008; Holt-Lunstad et al., 2010; Krämer, Sauer, and Ellison, 2021) in a way that maps onto the household-production decomposition; a richer specification with an additional shallow-tie input alongside depth would deliver the same equilibrium under the same calibration.

II. RELATED LITERATURE

The paper contributes to four literatures.

The fertility puzzle. Kearney, Levine, and Pardue (2022) document the post-2007 US fertility decline and show no economic, policy, or social factor commonly invoked can account for it. Kearney and Levine (2025) extend the puzzle to six high-income countries, formalize a “shifting priorities” framework, and conclude that period-price explanations cannot explain the widespread decline. We engage this conclusion directly: the hedonic phone-price decline is the supply-side shock, but its fertility consequences operate through a *cohort* channel — cohorts whose 18–25 partnership-formation window overlapped digital saturation accumulated lower depth stocks, with downstream consequences years later. Our cohort-exposure designs in Sections VI and VII make this empirically explicit. Doepke et al. (2023) survey the field and emphasize that the Beckerian quantity-quality framework no longer fits recent data; we extend it by applying the core logic to a new factor in household production. Mahler, Tertilt, and Yum (2025) model fertility decline as driven by social-comparison-induced intensive parenting, calibrated to Facebook-network data; their mechanism raises the cost of children rather than lowering the price of shallow connections.

Leisure technology. Aguiar et al. (2021) establish the leisure-luxury framework: improvements in recreational computing raised the marginal value of time, lowering young men’s labor supply. Yanagimoto (2024) extends this to the marriage and fertility margin with a leisure-bargaining model in which digital leisure is more valued by single individuals than partnered ones, because partnered consumption requires intra-household compromise. Calibrated to Japan, Yanagimoto’s model explains 21% of the marriage decline and 73% of the fertility decline. The mechanism is observationally similar to ours in aggregate prediction — both predict declining partnership formation as digital leisure becomes cheaper — but the two frameworks make distinct cross-sectional predictions about leisure consumption by partnership status. Our depth-investment framework predicts that partnered individuals consume less digital leisure but *more* in-person socializing

than otherwise-comparable unpartnered individuals; Yanagimoto’s bargaining framework predicts partnered individuals consume less of all leisure types. We test this discriminating prediction directly using ATUS time-use data and find the partnership-status coefficient on in-person socializing is positive, inconsistent with the bargaining channel as the primary mechanism.³

Marriage market frictions. Büyükeren et al. (2026) use Tinder’s rollout to identify causal effects on young adults; they find sharply increased sexual activity but no impact on long-term partnership formation, with rising dating-outcome inequality among men. This finding is consistent with the matching-margin mechanism in our model: digital matching technology can ease search frictions without producing more partnerships, because the partnerships that would form require depth investment that has been crowded out by digital time use.

Digital communication, social capital, and well-being. A substantial psychology literature documents the link between digital communication and social isolation (Twenge et al., 2018, 2019; Haidt, 2024; Turkle, 2011, 2015; Primack et al., 2017), with debate over magnitudes (Orben and Przybylski, 2019) but consistent direction. Section VI.A below imports causal evidence from our prior work on the same margin: applying a terrain-ruggedness IV on county-level digital infrastructure rollout yields a sharp decline in teen birth rates over 2003–2018, operating through reduced face-to-face socializing. An adjacent literature on loneliness and partnership formation (Cacioppo and Patrick, 2008;

³The test pools ATUS waves 2003–2024 and restricts to respondents 15 and older with non-missing partnership status ($n = 244,026$). Married respondents consume 33.16 fewer minutes per day of digital leisure than otherwise-comparable unpartnered respondents (cluster-robust SE 1.24, $t = -26.7$), consistent with both frameworks at the digital-leisure margin alone. Married respondents simultaneously consume 1.87 more minutes per day of in-person socializing than unpartnered respondents (cluster-robust SE 0.65, $t = 2.87$), with the coefficient growing monotonically from -1.6 in the 15–24 bracket (not significant) to $+3.6$ in the 45–54 bracket (significant at 1%) and crossing zero in the late twenties. The sign reversal between the two leisure types is the discriminating result: a leisure-bargaining mechanism would predict negative coefficients on both leisure types, while the depth-investment framework uniquely predicts the positive sign on in-person socializing alongside the negative sign on digital leisure, with the age gradient consistent with the depth-investment channel strengthening as within-partnership depth accumulates over the lifecycle. A complementary regression of completed-fertility-to-date on cohort broadband exposure during the respondent’s 18–25 window, restricted to partnered respondents 25–44 ($n = 57,379$), yields a coefficient of -0.385 (cluster-robust SE 0.025, $t = -15.3$), consistent with depth-stock erosion from broadband exposure during the relationship-formation years. All specifications include age-bracket fixed effects, year fixed effects, and standard demographic controls.

Holt-Lunstad et al., 2010; Murthy, 2023) establishes that loneliness is a consequence of perceived deficits in the *quality*, not the quantity, of social connections — the depth-versus-breadth distinction at the conceptual heart of our framework. The social-support literature (House, 1981; Cohen and Wills, 1985; Krämer, Sauer, and Ellison, 2021; Dunbar, 2018) provides direct empirical evidence for the asymmetric-substitution structure: instrumental support is delivered almost exclusively by strong ties and resists digital substitution; large weak-tie digital networks do not substitute for it. Putnam (2000) documents the secular decline in American social capital; Greenwood et al. (2005) provides the closest spiritual ancestor at the household-production level, with household appliances reshaping household production and labor force participation. Our contribution is to embed the depth-quality channel in a household-production framework with explicit fertility consequences.

III. A MODEL OF CONNECTION AND FERTILITY

III.A Two Becker tradeoffs

The model nests in the household-production tradition of Becker (1960); Becker and Lewis (1973); Becker (1981), in which children are chosen along quantity and quality margins with relative-price shifts driving substitution between the two; Galor and Weil (2000) closes this into a unified growth model in which industrialization endogenously delivers the first demographic transition. The framework’s limits have become visible in recent data: Doepke et al. (2023) observe that the regularities the framework was built to explain — the negative income-fertility gradient, the female labor force participation channel — no longer fit the post-2000 data cleanly, and the post-2007 acceleration into deep sub-replacement is not delivered by any standard parameterization (Kearney, Levine, and Pardue, 2022; Kearney and Levine, 2025).

We extend rather than abandon the Becker engine. Becker’s structural innovation was to recognize that household production has multiple inputs that respond to relative prices, and that fertility outcomes follow from how households reallocate across them. The first Becker tradeoff ran on children themselves — quantity versus quality of investment per

child. We argue a *second* tradeoff has emerged, running on a different input: the depth of social connections in which children are conceived and raised. Industrialization raised the return to child quality and triggered the first tradeoff; digitalization has lowered the relative price of breadth in connections and triggered the second. Digital tools sharply reduce the cost of maintaining many shallow ties without altering the cost of depth, and households reallocate accordingly.

The second tradeoff has fertility consequences only because depth enters the household production function for children as a complementary input. As depth erodes, the depth-to-birth function $b(D) = 1 - e^{-D/\kappa}$ slides off its saturation plateau into a region where birth probabilities respond strongly to D . Households that would have produced the Becker-optimal one or two children in a high-depth regime now have lower per-period birth probabilities; the first Becker tradeoff continues to operate only for the diminishing share of agents whose depth stock keeps them near saturation. Table 2 summarizes the parallel structure of the two tradeoffs.

TABLE 2. Two Becker Tradeoffs in Household Production

	<i>First tradeoff</i> (<i>Becker tradition</i>)	<i>Second tradeoff</i> (<i>this paper</i>)
Trigger	Industrial revolution	Digital revolution
Relative-price shift	Quality of children becomes cheaper relative to quantity (rising returns to human capital)	Breadth of connections becomes cheaper relative to depth (algorithmic broadcast and asynchronous tools)
Household substitution	Toward quality, away from quantity of children	Toward broad-shallow connections, away from deep ones
Fertility consequence	Quantity-quality reallocation: fewer, better-invested children	Depth-stock erosion makes the pre-condition for child-rearing scarce
Empirical signature	Demographic transition: TFR falls from 5–6 to 1.5–2.0	Sub-replacement acceleration: TFR falls from 1.5–2.0 toward 0.7–1.2

The paper is best read as a structural extension of unified growth theory: industrialization activated the first tradeoff and produced the demographic transition; digitalization has activated the second tradeoff and is producing the further collapse below replacement.

III.B Evidence the model must reproduce

The model is designed to reproduce four quantitative patterns in the cross-country panel: (i) the post-2007 break in period TFR, with magnitudes that vary across countries in proportion to the timing and depth of digital saturation; (ii) the rise in the share of cohorts reaching age 30 without a partnership, again proportional to digital saturation; (iii) the time-use shift from in-person socializing to digital leisure, with cross-country variation matching the same digital-saturation gradient; (iv) the conditional fertility decline among partnered agents, smaller in magnitude than the partnership-formation decline but contributing meaningfully to the aggregate.

The model abstracts from a number of forces that the literature has shown to be important elsewhere, including women’s wages, housing costs, contraceptive technology, and intensive-parenting demands. These are second-order for the post-2007 acceleration in our framework but can be reintroduced for robustness.

III.C The model

The economy reduces to a single fundamental tradeoff: agents allocate time between depth-building investment in close relationships and everything else. Each period t , an agent splits a unit of time three ways: depth investment h_t , market work ℓ_t , and outside-option time x_t :

$$h_t + \ell_t + x_t = 1. \tag{1}$$

The outside option x_t subsumes consumption, digital leisure, and the maintenance of broad-shallow connections — non-market activities that compete with depth investment. The breadth-versus-depth distinction central to the paper is preserved in the interpretation of x (which includes broad-shallow connection maintenance) versus h (which produces depth specifically). We do not model breadth as a separate stock because the predictions of interest are unchanged by collapsing breadth into the outside-option bundle. Market work ℓ_t is treated as exogenously fixed at $\bar{\ell}$ for each agent, set by labor-market institutions outside the model; this keeps focus on the depth-vs-leisure margin while letting wages

enter through the household budget.

The depth stock — which we also refer to as relational capital — evolves as

$$D_t = (1 - \delta)D_{t-1} + h_t, \quad (2)$$

with depreciation $\delta \in (0, 1)$. Depth is built only by direct time investment; nothing else produces it. This is the asymmetric-substitution structure of the paper, in its most parsimonious form: digital technology and other outside-option activities cannot substitute as inputs into D .

Fertility is a stochastic outcome of relational capital. In each period, an agent with depth D_t has a per-period probability of producing a birth equal to

$$b(D_t) = 1 - e^{-D_t/\kappa}, \quad (3)$$

where $\kappa > 0$ controls the curvature of the depth-to-birth mapping. The function $b(D)$ is increasing and saturating: low-depth agents have low birth probabilities, high-depth agents approach saturation, and marginal returns to depth diminish as the relationship becomes well-suited for child-rearing.

Per-period flow utility is

$$u_i(c_{i,t}, D_{i,t}; \gamma_i) = \log c_{i,t} + \gamma_i \log b(D_{i,t}), \quad (4)$$

where $c_{i,t}$ is the consumption-services flow household i enjoys during outside-option time, and $\gamma_i > 0$ is the household's preference for fertility. Households are heterogeneous in the preference parameter $\gamma_i \sim \text{Lognormal}(\mu_\gamma, \sigma_\gamma^2)$, drawn independently across households with the distribution fixed across time. Preferences are stationary: the function in (4) does not change over time.

What changes over time is the technology by which outside-option time is converted into consumption services. Outside-option activity combines time x_t with a market-purchased durable bundle (phones, the digital platforms they unlock, and complementary

services) at hedonic-quality-adjusted price p_t , financed by labor income $w_t \bar{\ell}$ from a fixed labor allocation $\bar{\ell}$. The exogenous price-and-wage component of the quality bundle is $Q_t = w_t \bar{\ell} / p_t$.

Crucially, the consumption-services value of the digital bundle depends on the bundle's network of users, not only on its hedonic price. A phone is more valuable to its owner when more peers are on the same platforms: more group chats, more posts to react to, more content tuned to the user's social and cultural niche, more friends to coordinate plans with. We capture this with a network multiplier $\psi(A_t^c)$ on the quality bundle, where $A_t^c \in [0, 1]$ is the population-mean leisure share within the agent's reference population c — the share of non-depth time the relevant peer group allocates to the digital outside option. The agent's effective quality bundle is then

$$\tilde{Q}_t = Q_t \cdot \psi(A_t^c), \quad \psi(A) = 1 + \frac{\psi_{\max}}{1 + \exp(-k(A - A^*))}, \quad (5)$$

with $\psi_{\max} > 0$ the saturation value of the network amplification, $A^* \in (0, 1)$ the inflection point, and $k > 0$ the logistic steepness. The reference population c is country-specific: language and cultural barriers mean that the relevant peer network for an Italian household is largely composed of other Italians, and the network value of an Italian's Instagram session depends on Italian aggregate engagement, not on the global aggregate.⁴

The teen-age special case. The network multiplier $\psi(A_t^c)$ is the formal counterpart of the coordination-collapse mechanism we develop in Section VI.A, where the post-2007 collapse of teen in-person social activity is shown to be too coordinated across teens to be explained by independent optimization over individual price changes, and is instrumentally tied to county-level digital infrastructure rollout. Equation (5) formalizes this: as more individuals shift toward the digital outside option, A_t^c rises, $\psi(A_t^c)$ rises, the productivity of the digital bundle for any individual rises, and the shift reinforces itself. In the limit

⁴Under an alternative specification with a single global A_t common to all countries, the substantive results below are essentially unchanged (Section V.F). The country-specific specification is theoretically more defensible. We treat A_t^c as the population-mean within country c and treat each country as solving its own equilibrium fixed point.

where the depth stock has not yet accumulated and the unintended-conception margin dominates the chosen-fertility margin (the youngest-age cohort), the household problem below collapses to a contemporaneous time-allocation problem in which $\psi(A_t^c)$ is the dominant force. The present paper generalizes that result across the full partnership-formation window (15–29): same supply-side shock, same country-specific equilibrium feedback, same network amplifier, with depth-stock dynamics activating as cohorts age in.

We parameterize the household-production function for consumption services as a CES aggregator of time and the network-amplified quality bundle, with elasticity of substitution σ and quality share weight α :

$$c_{i,t} = [\alpha \tilde{Q}_t^{(\sigma-1)/\sigma} + (1-\alpha)x_{i,t}^{(\sigma-1)/\sigma}]^{\sigma/(\sigma-1)}, \quad (6)$$

with $\sigma \in (0, 1)$ implying *gross complements* between time and quality and $\alpha \in (0, 1)$ governing the relative weight of quality. The household-specific input is leisure time $x_{i,t}$; the effective quality bundle $\tilde{Q}_t = Q_t \psi(A_t^c)$ is common to all households in the country. The gross-complements assumption captures the substantive economics: phones make outside-option time more valuable, not less. When the quality bundle becomes cheaper (rising w_t/p_t) or the network larger (rising A_t^c), a household with the same time allocation enjoys a richer consumption flow, and the marginal value of additional outside-option time *rises* because each minute now combines with more quality.⁵

Substituting (6) into (4) and using $x_{i,t} = 1 - h_{i,t}$ (with the labor allocation $\bar{\ell}$ absorbed into the calibration of δ and κ as a normalization), the per-period utility evaluated at chosen $h_{i,t}$ is

$$u_{i,t} = \frac{\sigma}{\sigma-1} \log[\alpha \tilde{Q}_t^{(\sigma-1)/\sigma} + (1-\alpha)(1-h_{i,t})^{(\sigma-1)/\sigma}] + \gamma_i \log b(D_{i,t}). \quad (7)$$

⁵The Cobb-Douglas case $\sigma = 1$ collapses (6) to $c_{i,t} = \tilde{Q}_t^\alpha x_{i,t}^{1-\alpha}$, which substituted into log utility gives $u_{i,t} = \alpha \log \tilde{Q}_t + (1-\alpha) \log x_{i,t} + \gamma_i \log b(D_{i,t})$. The first term is additive in time and drops out of the FOC, so the depth-vs-leisure margin would be invariant to \tilde{Q}_t . This is a textbook result: with logarithmic utility and Cobb-Douglas household production, income and substitution effects exactly cancel. Gross complements ($\sigma < 1$) break this knife-edge: rising \tilde{Q}_t raises the marginal product of $x_{i,t}$ in producing $c_{i,t}$, so the substitution effect dominates and the agent reallocates time from depth to leisure.

The marginal utility of $x_{i,t}$ is $\partial \log c_{i,t} / \partial x_{i,t} = \omega_{i,t} / x_{i,t}$, where

$$\omega_{i,t} \equiv s(\tilde{Q}_t, x_{i,t}) = \frac{(1 - \alpha) x_{i,t}^{(\sigma-1)/\sigma}}{\alpha \tilde{Q}_t^{(\sigma-1)/\sigma} + (1 - \alpha) x_{i,t}^{(\sigma-1)/\sigma}} \quad (8)$$

is the time share of consumption services in the CES bundle. This is the productivity of outside-option time in the same form as the standard one-input log-utility case: with stationary preferences over $(c_{i,t}, D_{i,t})$, the agent's depth-vs-leisure FOC is governed by the time share $\omega_{i,t}$, which under (8) is itself an endogenous function of the effective relative price $\tilde{Q}_t = Q_t \psi(A_t^c)$ and the equilibrium $x_{i,t}$. Under gross complements ($\sigma < 1$), $\omega_{i,t}$ is increasing in \tilde{Q}_t : cheaper quality *or* a larger network raises the time share, raising the marginal value of outside-option time and shifting the depth-vs-leisure margin away from depth.

The household discounts the future at rate $\beta \in (0, 1)$. Given a deterministic exogenous path $\{Q_t\}_{t \geq 0}$ for the country, an aggregate-network path $\{A_t^c\}_{t \geq 0}$ taken as given by the individual agent, household-specific preference γ_i , and initial depth stock $D_{i,-1}$, the value function $V_{i,t}(D_{i,t-1}; \gamma_i)$ satisfies the Bellman equation

$$V_{i,t}(D_{i,t-1}; \gamma_i) = \max_{h_{i,t} \in [0,1]} \left\{ \log c_{i,t}(\tilde{Q}_t, 1 - h_{i,t}) + \gamma_i \log b(D_{i,t}) \right\} + \beta V_{i,t+1}(D_{i,t}; \gamma_i), \quad (9)$$

with $D_{i,t} = (1 - \delta)D_{i,t-1} + h_{i,t}$, $\tilde{Q}_t = Q_t \psi(A_t^c)$, $c_{i,t}$ given by (6), and the maximization understood as joint over flow and continuation since $D_{i,t}$ depends on $h_{i,t}$.

The interior first-order condition with respect to $h_{i,t}$ is

$$\underbrace{\frac{\omega_{i,t}}{1 - h_{i,t}^*}}_{\text{marginal cost}} = \underbrace{\gamma_i \cdot \frac{b'(D_{i,t}^*)}{b(D_{i,t}^*)}}_{\text{flow benefit}} + \underbrace{\beta \frac{\partial V_{i,t+1}}{\partial D_{i,t}} \Big|_{D_{i,t}=D_{i,t}^*}}_{\text{continuation value}}, \quad (10)$$

where $\omega_{i,t}$ is the equilibrium time share from (8) and $b'(D) = e^{-D/\kappa} / \kappa$. The envelope

condition for the future value is

$$\frac{\partial V_{i,t+1}}{\partial D_{i,t}} = (1 - \delta) \cdot \frac{\omega_{i,t+1}}{1 - h_{i,t+1}^*}, \quad (11)$$

which substituted into (10) and rolled forward gives the recursive characterization

$$\frac{\omega_{i,t}}{1 - h_{i,t}^*} = \gamma_i \frac{b'(D_{i,t}^*)}{b(D_{i,t}^*)} + \beta(1 - \delta) \frac{\omega_{i,t+1}}{1 - h_{i,t+1}^*}. \quad (12)$$

Equation (12) is the depth-investment Euler equation: at the optimum, today's marginal cost of depth investment equals the flow marginal benefit plus the discounted, depreciation-adjusted marginal value of depth tomorrow. The economic content of $\omega_{i,t}/(1 - h_{i,t})$ is the marginal opportunity cost of depth investment: each additional unit of $h_{i,t}$ removes a unit from $x_{i,t}$, reducing consumption services by $\omega_{i,t}/(1 - h_{i,t})$. When p_t falls or w_t rises, Q_t rises; when A_t^c rises, $\psi(A_t^c)$ rises; either channel raises \tilde{Q}_t , the time share $\omega_{i,t}$, and the opportunity cost of depth investment — not because the agent has changed preferences, but because each minute spent on depth-building forgoes a richer consumption bundle. Note that the heterogeneous γ_i on the right-hand side of (12) implies that high- γ_i households tolerate a higher marginal opportunity cost before reducing depth investment, while low- γ_i households contract $h_{i,t}^*$ further toward the corner; the model's cross-household variation in fertility outcomes is generated entirely by the dispersion of γ_i .

Equilibrium definition. The aggregate A_t^c entering each agent's problem is the population-mean leisure share:

$$A_t^c = \mathbb{E}_i[1 - h_{i,t}^*(A_t^c; \gamma_i)], \quad (13)$$

where the expectation is over the heterogeneity distribution of γ_i within country c and the policy $h_{i,t}^*$ depends on both the household's own γ_i and the aggregate trajectory $\{A_s^c\}_{s \geq 0}$ that the agent treats as exogenous. An equilibrium is a trajectory $\{A_t^c\}$ such that the population mean of optimal individual policies reproduces the trajectory the agents took as given. Each country solves its own equilibrium fixed point with its own $\{Q_t^c\}$ path;

the household primitives $(\delta, \kappa, \sigma_\gamma, \beta, \mu_\gamma, \alpha, \sigma, \psi_{\max}, A^*, k)$ are common across countries, while γ_i varies across households within each country.

The system is solved by Krusell–Smith iteration on the trajectory $\{A_t^c\}$. Inner step: given a guess for $\{A_t^c\}$, every agent’s problem reduces to a deterministic dynamic problem with effective quality path $\{Q_t^c \psi(A_t^c)\}$, solved by the endogenous-grid method (Carroll, 2006) — given the policy-target $D_{i,t}^*$, the FOC determines $h_{i,t}^*$ via a one-dimensional rootfind on the nonlinear LHS $\omega_{i,t}(\tilde{Q}_t, 1 - h_{i,t})/(1 - h_{i,t})$, which is monotone increasing in $h_{i,t}$. Outer step: aggregate the resulting policies via (13) to obtain \hat{A}_t^c and update the guess. We use damping $\theta = 0.5$ and convergence tolerance $\|\hat{A}^c - A^c\|_\infty < 10^{-3}$. Convergence in our calibrated parameter region is monotone and takes 10–15 iterations; we numerically verify uniqueness of the equilibrium trajectory by perturbing the initial guess.

Country heterogeneity through wages and networks. The model has two sources of cross-country variation, both grounded in (5)–(8). First, the hedonic price p_t is determined by a globally integrated supply chain — phones manufactured for the world market are roughly the same physical goods at the same quality-adjusted prices everywhere (Byrne, 2019) — but wages w_t are country-specific. Korea’s real wages rose 32 percent over 2003–2023; Italy’s fell 3 percent. Different w_t/p_t paths translate into different Q_t^c paths. Second, the network multiplier $\psi(A_t^c)$ depends on each country’s own population aggregate A_t^c , which is itself the equilibrium response of the country’s population to its own $\{Q_t^c\}$. Countries with steeper Q_t^c rises generate higher A_t^c at any given t , which in turn raises $\psi(A_t^c)$. Both channels operate jointly in the cross-country exercise (Section V.G); we will see that in the calibrated parameter region the wage channel does essentially all the heavy lifting, with the network channel set-identified at a configuration that produces negligible amplification within the empirical A_t^c range.

In the calibration of Section IV, we solve the full dynamic problem (9) numerically: backward induction on a state grid for D , with the endogenous-grid method (Carroll, 2006) delivering the optimal policy at each (t, D, γ_i) in vectorized form, followed by forward simulation from the backcast initial condition. The discount factor is set to $\beta = 0.96$ per

year, the standard annual choice corresponding to a four-percent time discount rate.

III.D Comparative statics

The four propositions below follow from the FOC.

Proposition 1 (Time reallocation). h_t^* is decreasing in ω_t . As digital technology raises the productivity of outside-option time, agents substitute time away from depth investment.

Proposition 2 (Depth-stock erosion). The steady-state depth stock $\bar{D}^* = h^* / \delta$ is decreasing in ω . Along any transition path induced by a rising $\{\omega_t\}$, the depth stock falls monotonically toward its new steady state.

Proposition 3 (Fertility decline). The aggregate per-period birth rate $b(D^*)$ is decreasing in ω . As ω rises and D^* falls, the depth-to-birth function maps directly onto a falling birth rate without any threshold non-linearity required.

Proposition 4 (Heterogeneous response). The response to a given $\omega_{i,t}$ shock is larger for low- γ_i households than for high- γ_i households: at low γ_i the marginal benefit of depth is small relative to the marginal opportunity cost, so $h_{i,t}^*$ falls toward zero; at high γ_i the marginal benefit remains large enough to maintain meaningful $h_{i,t}^*$ even at elevated $\omega_{i,t}$. Population aggregate fertility falls through both the intensive-margin contraction of birth probabilities at all γ_i and the extensive-margin movement of low- γ_i households toward the corner.

The model's central qualitative prediction is that rising $\{\omega_t\}$ — driven by the falling hedonic price $\{p_t\}$ and the rising network multiplier $\{\psi(A_t^c)\}$ — lowers all three: time invested in depth, the depth stock, and the resulting birth rate. These are not independent margins; they are sequential consequences of the same shock. The quantitative bite depends on the calibrated values of δ (how fast depth depreciates), κ (how steeply depth maps into births), σ_γ (how heterogeneous the population is), σ and α (CES curvature and quality-share weight), and ψ_{\max} and A^* (saturation amplification and inflection point of the network multiplier). In Sections IV and V we calibrate these parameters on US 2007–2024

moments and show the implied trajectory accounts for essentially the full US 2007–2024 fertility decline and captures the post-2007 OECD aggregate decline through 2023.

III.E Mapping to data

The model produces a time path of population-mean depth \bar{D}_t and a time path of population-mean birth probability \bar{b}_t . We map birth probability to the empirical K_t series (children per adult age 25–44 in ATUS) through a units-mapping constant: $K_t^{\text{model}} = c_K \cdot \bar{b}_t$. Children appear in the data as a stock (children currently in the household) but in the model as a flow (per-period birth probability); the units constant is calibrated by ordinary least squares to match the observed K_t levels. This treats the model’s \bar{b}_t as a steady-state proxy for the implied stock, which is exact when the demographic rates are stationary and a reasonable approximation along slowly-evolving transition paths such as the one we calibrate.

IV. CALIBRATION

We implement the model on US data over the 2007–2024 smartphone era. The empirical moments are population-mean share of discretionary time on digital outside-option activities (computed from ATUS) and children per adult age 25–44 (computed from ATUS partnership-status data in Section I.A). The calibration window starts in 2007, the launch of the iPhone and the empirical onset of the post-2007 fertility break documented in Section I.A. This is the simplest setting in which the model can be tested against the empirical phenomenon it is designed to explain; cross-country calibration with country-specific digital-saturation timing remains as the natural extension.

IV.A Data: hedonic prices and real wages

The empirical moments are computed directly from the BLS American Time Use Survey activity-summary microdata (atussum 0324, $n = 252,808$ respondent-day observations), weighted with TUFNWGTP. Time on digital outside-option activities is the sum of ATUS codes for television (120303, 120304), playing games (120307), and computer use

for leisure (120308). This definition is broader than the “digital leisure” series of Fact 2 in Section I.A, which excludes television. The empirical K_t series is the population-weighted mean number of household children under 18 per adult aged 25–44.

The central identification choice is how to pin down the time path of $\{\tilde{Q}_t\}$, the network-amplified effective quality bundle. Recall from (5) that $\tilde{Q}_t = Q_t \psi(A_t^c)$, with $Q_t = (w_t/p_t)/(w_{2003}/p_{2003})$ the wage-to-price ratio normalized to 2003 (an exogenous time series for each country) and $\psi(A_t^c)$ the network multiplier (an equilibrium object determined within the country’s Krusell–Smith fixed point). The price series $\{p_t\}$ is the same global supply-side shock for all countries (a globally integrated supply chain); the wage series $\{w_t\}$ is country-specific; the network multiplier varies across countries through their equilibrium A_t^c trajectories.

Hedonic price series. Our primary phone-price series is the Aizcorbe et al. (2019) “Overall mobile phones” Fisher hedonic index (henceforth ABS-2019), constructed by Federal Reserve Board economists from IDC retail data using time-varying hedonic regressions on processor speed, RAM, screen size, camera resolution, manufacturer, operating system, and mobile generation. ABS-2019 has been adopted by the Bureau of Economic Analysis for the GDP accounts since 2018 and is the same series we use in Section VI.A. The index is published for 2010–2017 (their Table 2, line 7) with annual percentage changes $-9.6, -14.9, -15.2, -28.0, -18.8, -22.5, -15.3$, averaging -17.8 percent per year. Because ABS-2019 covers only 2010–2017 we extend it: chained backward through the BLS hedonic CPI for telephone hardware (CUUR0000SEEE01) for 1997–2009 and forward through the same series for 2018–2024, with the link points at 2010 and 2017 respectively. After deflating by all-items CPI, the resulting ABS-extended series fell from 5.06 in 1997 to 1.00 in 2003 to 0.0140 in 2024 — a 98.6 percent real decline 2003–2024.

Real wage series. Real wage data come from the OECD “Average Annual Wages” series, expressed in 2024 constant US dollars at PPP, available 1995–2024 for all 38 OECD

member countries. US real wages rose from \$66,662 in 2003 to \$82,933 in 2024, a 24 percent cumulative gain. Other countries in the OECD panel show different trajectories: CEE members roughly doubled (Lithuania +108 percent, Latvia +95 percent, Estonia +74 percent) reflecting post-accession convergence; Korea +32 percent and Norway +34 percent; major Anglo-Saxon economies grew 10–25 percent; Mediterranean economies flat to negative (Italy –3 percent, Greece –15 percent). The cross-country dispersion in wage growth is substantial and is the main source of cross-country variation in the model; we use the full 38-country panel in the cross-country exercise (Section V.G).

Effective relative price and \tilde{Q}_t . Combining the hedonic price and the wage gives $Q_t = (w_t/p_t)/(w_{2003}/p_{2003})$, normalized to 2003. For the US, $Q_{2024} \approx 88.6$: the wage-deflated effective price (the inverse p_t/w_t normalized to 2003 = 1) fell from 1.00 to 0.0113 — a 98.9 percent decline — with rising wages compounding with the falling price. Multiplying by the equilibrium network multiplier $\psi(A_{2024}^{US}) \approx 1.02$ gives the network-amplified effective quality bundle $\tilde{Q}_{2024} \approx 90$; the network amplifier contributes essentially zero amplification at the calibrated parameters. This effective relative price is what enters ω_t via (8). The price-and-wage shock enters the household problem through the household-production technology (eq. 6), not through preferences (eq. 4); preferences over the consumption-fertility bundle (c, D) are stationary throughout.

IV.B Initial condition: backcast from 1993

Rather than imposing an arbitrary scaling on the initial depth-stock, we run the model from a 1997-conditions steady state initialized in 1993 and let the simulation determine the state entering the calibration window. We assume the effective relative price $p_t/w_t = (p/w)_{1997}$ for all $t < 1997$, putting the system in steady state at $\tilde{Q}_{1997} = Q_{1997} \psi(A_{1997}^c)$ at the start of 1993, then run the actual trajectory $\{Q_t\}_{t=1997}^{2024}$ forward, where $Q_t = (w_t/p_t)/(w_{2003}/p_{2003})$ is the wage-to-price ratio normalized to 2003 and A_t^c is determined in equilibrium. This eliminates a free initial-condition parameter. Figure 8 plots the extended price series and the implied equilibrium time-share ω_t trajectory under

the calibrated network multiplier.




FIGURE 8. From hedonic phone price to outside-option time share, with pre-1997 backcast. Left panel: real relative phone price $\rho_t = p_t/p_{2003}$ (ABS-2019 extended via BLS CUUR0000SEEE01, deflated by CPI-U, 2003 = 1). Pre-1997 held at the 1997 level (gray segment). Right panel: equilibrium time share $\omega_t = s(\tilde{Q}_t, x_t^*)$ implied by the calibrated CES technology under US wages and equilibrium A_t^c network amplification.

IV.C Parameters and identification

Five structural parameters are fixed from external evidence. The depth depreciation rate is set to $\delta = 0.15$ per year. Roberts and Dunbar (2011) track 25 close ties over 18 months without contact and find emotional closeness declines roughly linearly at 15–20 percent per year for inner-circle friends; marital satisfaction in the absence of investment declines at 5–10 percent per year (Bradbury and Karney, 2004). The aggregate depth-decay parameter sits between these. The depth-to-birth conversion parameter is fixed at $\kappa = 10$, anchored to median-preference households producing approximately replacement-level fertility in pre-shock steady state. The log-standard-deviation of the preference distribution is fixed at $\sigma_\gamma = 0.55$, anchored to the cross-sectional standard deviation of completed cohort fertility for US women born 1955–1965 (mean ≈ 2.0 , sd ≈ 1.2 , log-sd ≈ 0.55) — the pre-digital cohort that completed fertility before our window. The discount factor is set to $\beta = 0.96$, the standard annual choice corresponding to a 4 percent time discount rate. The logistic steepness in the network multiplier (5) is fixed at $k = 20$, a default that gives a transition width of about $1/k = 0.05$ in A units; sensitivity to k is reported in Section V.F. Fixing σ_γ , β , and k externally is necessary because the time-series moments cannot pin

them down separately from $(\alpha, \sigma, \psi_{\max}, A^*)$.

We calibrate five parameters jointly: $(\mu_\gamma, \alpha, \sigma, \psi_{\max}, A^*)$. Households are heterogeneous in $\gamma_i \sim \text{Lognormal}(\mu_\gamma, \sigma_\gamma^2)$ with the distribution fixed across time. For each γ_i in a 21-point Gauss-Hermite grid, we solve the household’s dynamic problem (9) by backward induction over a 60-point log-spaced grid for $D_{i,t} \in [0.01, 12]$ and a 1993–2100 horizon (the terminal date is well beyond the 2007–2024 calibration window). The optimal policy $h_{i,t}^*$ is computed by the endogenous-grid method (Carroll, 2006) via vectorized bisection in $h_{i,t}$ on a monotone FOC. Forward simulation starts from the dynamic stationary state at \tilde{Q}_{1997} . The aggregate A_t^c entering each agent’s problem is solved by Krusell–Smith outer-loop iteration. Calibration targets six moments: \bar{x}_t at 2007, 2010, 2017, 2024, and \bar{K}_t at 2007 and 2024.

Identification is broadly clean. σ governs the curvature of the CES and is pinned by the slope of \bar{x}_t in the early portion of the window when A_t^c is far below A^* and the network multiplier is essentially inactive. α governs the level of ω_t at the 2007 baseline and is pinned by the level of \bar{x}_t in 2007. μ_γ governs the level of \bar{K}_t and is pinned by the 2007 level. ψ_{\max} governs the asymptotic amplification of the digital bundle and is pinned by the level of \bar{x}_t in 2024. A^* governs the timing of the transition from the network-inactive regime to the network-active regime and is pinned by the shape of the \bar{x}_t trajectory in the middle of the window (the 2010 and 2017 values).

One feature of the calibrated solution is worth flagging. The equilibrium A_t^c trajectory rises from approximately 0.25 in 1997 to 0.53 in 2024, peaking well below the calibrated inflection point $A^* = 0.715$. The network multiplier therefore operates throughout the calibration window in the deep sub-inflection regime of the logistic — amplification rises smoothly from $\psi(A_{1997}) \approx 1.00$ to $\psi(A_{2024}) \approx 1.02$, contributing essentially zero amplification on top of the wage-deflated Q_t channel. Both ψ_{\max} and A^* are set-identified rather than point-identified at the calibrated values; the relevant empirical conclusion is that the calibrated moments do not require any meaningful network amplification on top of the underlying CES-with- Q_t channel. The network specification reduces in practice to the no-network CES specification we report as a robustness check (Section V.F); the two

produce essentially identical headline numbers.

Table 3 reports the calibrated values: $\sigma = 0.53$ (gross complements, well below 1), $\alpha = 0.64$, $\mu_\gamma = 0.28$ (median $\gamma_i = 1.32$), $\psi_{\max} = 0.709$, $A^* = 0.715$. The implied γ_i distribution has median 1.32 and P10/P90 range 0.65–2.67.

TABLE 3. Calibrated and externally-fixed parameter values

Parameter	Value	Source / interpretation
δ	0.150	Fixed (close-tie depreciation, Roberts & Dunbar 2011)
κ	10.0	Fixed (depth-to-birth curvature)
σ_γ	0.55	Fixed (completed-fertility CV from pre-digital US cohort)
β	0.96	Fixed (annual time discount, 4% rate)
k	20	Fixed (logistic steepness in network multiplier)
μ_γ	0.277	Calibrated \rightarrow median $\gamma_i = 1.32$
α	0.644	Calibrated CES quality-share weight
σ	0.530	Calibrated CES elasticity of substitution (gross complements)
ψ_{\max}	0.709	Calibrated saturation network amplification (set-identified)
A^*	0.715	Calibrated network inflection point (set-identified)
c_K	3.93	Units mapping (births per period to ATUS-stock units)

Initial condition: dynamic stationary state at \tilde{Q}_{1997} , then forward through actual \tilde{Q}_t .
Implied Q_t trajectory (US): 0.18 (1997) \rightarrow 1.00 (2003) \rightarrow 3.24 (2007) \rightarrow 88.6 (2024)
Implied A_t^c trajectory (US): 0.25 (1997) \rightarrow 0.41 (2003) \rightarrow 0.49 (2007) \rightarrow 0.53 (2024)
Implied $\psi(A_t^c)$ trajectory (US): 1.000 (1997) \rightarrow 1.002 (2003) \rightarrow 1.008 (2007) \rightarrow 1.017 (2024)

V. RESULTS

V.A Three joint predictions

The model’s central claim is that digital technology erodes relational capital, the depth-stock D that households build through time investment in close relationships. As D falls, three demographic outcomes follow jointly: lower marriage rates (because D is the input to forming durable partnerships), lower divorce rates conditional on marriage (because the surviving partnerships are self-selected toward high-commitment types), and lower fertility (because D is the input to producing children). Figure 9 verifies these against the calibrated model. Panel (A) shows population-mean relational capital \bar{D}_t declining from 4.06 in 2007 to 3.18 in 2024 (–21.6 percent). Panel (B) reports the share of households

above an active-fertility threshold ($\mu(D_t) \geq 0.15$, smoothed via a sigmoid with bandwidth 0.015 to avoid quadrature-discreteness artifacts) falling over the window: directionally consistent with the empirical 45 percent decline in 18–29 first-marriage rates (Figure 6), though smaller in magnitude. Panel (C) shows the selection ratio $E[\gamma_i | \text{active}]/E[\gamma_i]$ rising: the directional analogue of the falling divorce rate among surviving partnerships (Figure 5).




FIGURE 9. Three joint predictions of the model. Panel (A): population-mean relational capital \bar{D}_t falls from 4.06 in 2007 to 3.18 in 2024 (−21.6 percent). Panel (B): share of the population above the active-fertility threshold ($\mu(D) \geq 0.15$, sigmoid-smoothed) falls over the calibration window. Panel (C): selection ratio $E[\gamma_i | \text{active}]/E[\gamma_i]$ rises over the window.

The simultaneous decline of marriage and divorce is itself diagnostic for the relational-capital channel. Rising childcare costs reduce conditional fertility but have no clear implication for divorce stability. Career-versus-family conflict tends to lower marriage but raise divorce, the opposite of what the data show. Only a mechanism in which the same input supports partnership entry and maintenance — and becomes scarce on both simultaneously — produces the joint pattern: depth is required to enter a partnership and to sustain one; as depth becomes scarce, fewer households cross the entry threshold, and those who do are self-selected toward depth-rich types whose marriages survive at higher rates. Quantitatively, the model captures the full fertility decline through a combination of partnership-formation decline (the extensive margin) and conditional-fertility decline (the intensive margin), but it captures only a fraction of the marriage-rate and divorce-rate movements; the residual would require explicit partnership-formation frictions beyond the scope of this paper.

V.B Baseline calibration fit

Figure 10 plots the model-implied trajectories of \bar{x}_t and \bar{K}_t against their empirical counterparts. With five free parameters ($\mu_\gamma, \alpha, \sigma, \psi_{\max}, A^*$) identified by six moments (\bar{x}_t at 2007, 2010, 2017, 2024 and \bar{K}_t at 2007, 2024), the calibration achieves a tight fit at all six target moments. Each moment is matched within ± 0.4 percent. The empirical \bar{K}_t decline from 1.305 in 2007 to 1.064 in 2024 is matched by a model-implied decline from 1.306 to 1.062, accounting for essentially the entire observed change (fit share 101 percent). The empirical \bar{x}_t trajectory — 0.492 in 2007, 0.523 in 2010, 0.539 in 2017, 0.528 in 2024 — is matched by 0.492, 0.522, 0.537, 0.530. The shape of \bar{x}_t , which is concave (rising rapidly through the late 2000s and then plateauing), is captured by the equilibrium A_t^c trajectory rising toward but not crossing A^* ; the calibrated $\psi(A_t^c)$ trajectory contributes negligible amplification on top of the underlying CES-with- Q_t channel (rising from 1.000 in 1997 to 1.017 by 2024).

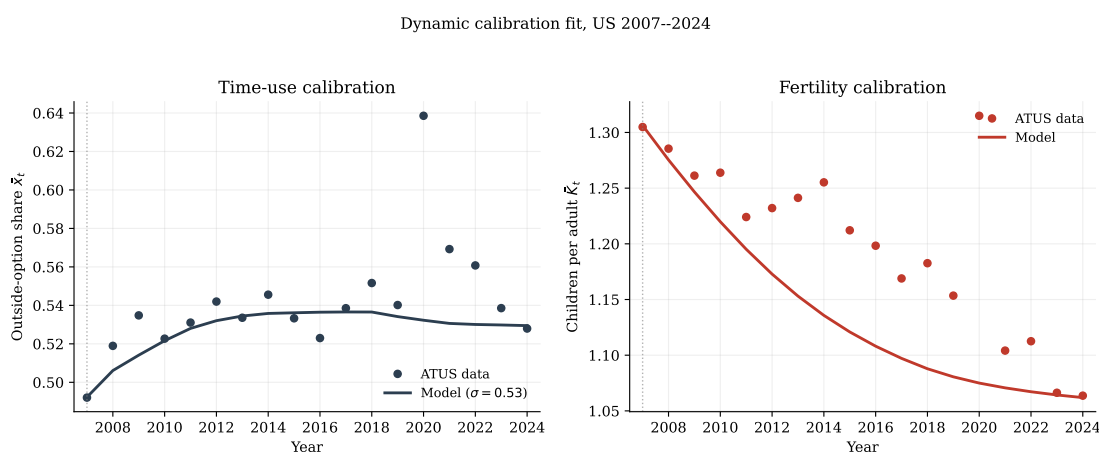


FIGURE 10. Model fit, US 2007–2024. Left panel: outside-option time share \bar{x}_t in the model and ATUS data. Right panel: children per adult age 25–44 in the model and ATUS data. The network-CES calibration matches all six target moments tightly (within ± 0.4 percent). The within-window front-loading of the model trajectory relative to the back-loaded data is visible in the \bar{K}_t panel.

Within-country year-by-year shape. The model matches the cumulative 2007–2024 decline but front-loads it relative to the data. This reflects perfect-foresight optimization

against a wage-deflated phone price that falls fastest in the early window; an imperfect-foresight extension would produce a more back-loaded trajectory.⁶

A second consideration suggests that the model’s single-shock specification, rather than perfect foresight itself, is the binding source of the front-loading. The digital regime has two distinct technological channels operating on different schedules. The device-side shock — captured by the hedonic phone price ρ_t in our model — is largely concluded by 2014, with the steepest real-price declines occurring between 2007 and 2012. The platform-side channel — attention-extraction intensity, proxied by Meta’s worldwide average revenue per user — only began rising from a near-zero base in 2011 (\$4.34) and grew roughly tenfold through 2023 (\$44.60), with the ad-tech infrastructure that drives this growth (News Feed ads in 2012, mobile ads in 2012, Custom Audiences in 2013, dynamic product ads in 2014, programmatic optimization through 2018) maturing entirely after the iPhone launch.⁷ The two channels are offset by roughly 5–7 years. A richer specification with both shocks on their realized schedules would generate a less front-loaded model trajectory even under perfect foresight, since households would not see the full platform-side intensification at $t = 0$ if it had not yet been built.

V.C Cumulative magnitude: the smartphone-era digital trajectory

The total-decline fit reported in Section V.B (101 percent of the US 2007–2024 decline) is a cumulative-magnitude claim, not a claim about the precise timing of the within-window decline. The model’s trajectory — driven by the wage-deflated phone-price decline from $Q_{2007} = 3.24$ to $Q_{2024} = 88.6$ — produces a monotone fertility decline whose cumulative magnitude through 2024 matches the data exactly, with the front-loaded shape discussed in the previous subsection. This is the claim the calibration supports and that the cross-country and aggregate exercises extend.

⁶The front-loading is structural to perfect-foresight optimization rather than a feature of any specific parameter choice. We verified this across multiple alternative specifications — varying network steepness k , the inflection A^* , the discount factor β , depth depreciation δ , replacing the dynastic Bellman with a 30-period OLG cohort structure, adding age-fecundability weights (Eijkemans et al., 2014; Steiner and Jukic, 2016), and using an age-specific peer network — and none changes the within-window shape materially.

⁷Meta investor disclosures; pre-2011 ARPU is not meaningfully reportable because Facebook had no significant ad infrastructure prior to that date and reported total revenue was small relative to its user base.

The within-window front-loading is visible when the OECD-aggregate fit is decomposed across windows. OECD-aggregate TFR declined 10.8 percent from 2007 to 2019 and 18.4 percent from 2007 to 2023. The US-calibrated model (applied with the same hedonic price and US wages, the relevant aggregate-identification stance discussed in Section V.H below) predicts -17.3 percent from 2007 to 2019 (a 159 percent overshoot of the shorter window) and -18.5 percent from 2007 to 2023 (a 101 percent fit of the longer window). The cumulative magnitude through 2023 matches the OECD aggregate exactly; the 2007–2019 sub-window does not because of the model’s front-loading. This is the same shape mismatch made visible at the aggregate level: cumulative through 2024 lines up; short windows do not.

The honest framing is that the model captures the cumulative magnitude of the post-2007 fertility movement through 2024 in both the US calibration target and the OECD aggregate falsification, and the calibration supports digital technology as a sufficient candidate explanation for the magnitudes observed over the smartphone-era window. The model neither explains nor claims to explain the precise within-window timing of the break or the pre-2007 trajectory of the high-income aggregate (which was rising).

V.D Counterfactual: phone price frozen at 2007

A second way to quantify the digital channel’s contribution is the standard model-internal counterfactual. We hold the effective relative price p_t/w_t at its 2007 value for all subsequent years (so Q_t remains at $Q_{2007} = 3.24$) and re-solve the equilibrium $\{A_t^c\}$ with this counterfactual price path, leaving the pre-2007 portion at its actual values. The counterfactual equilibrium has A_t^c peaking lower than the baseline equilibrium; the network amplifier is essentially inactive in both equilibria so only the Q_t channel responds. Figure 11 reports the result. The counterfactual \bar{K}_t trajectory declines from 1.306 in 2007 to 1.167 in 2024, against the baseline path’s 1.062. The cumulative gap between actual and counterfactual by 2024 is 0.105 children per adult age 25–44, which is **43 percent of the observed 2007–2024 decline of 0.241 children per adult**.

The 43 percent counterfactual share is the model’s estimate of the post-2007 incre-

mental movement in digital quality contribution to the post-2007 fertility decline. It is bounded below the 101 percent total fit because households respond to the entire $\{Q_t\}$ path, including its pre-2007 portion; the pre-2007 segment of the shock contributes a further share of the within-window decline that is allocated to “pre-2007 momentum” under this counterfactual definition (the depth stock continues to erode after the price-shock plateau because $\delta = 0.15$ implies a five-year depth half-life). Read as a marginal contribution statement, the 43 percent figure isolates how much additional fertility decline is produced by continued post-2007 movement in Q_t above and beyond what would have occurred had Q_t stabilized at its 2007 level.

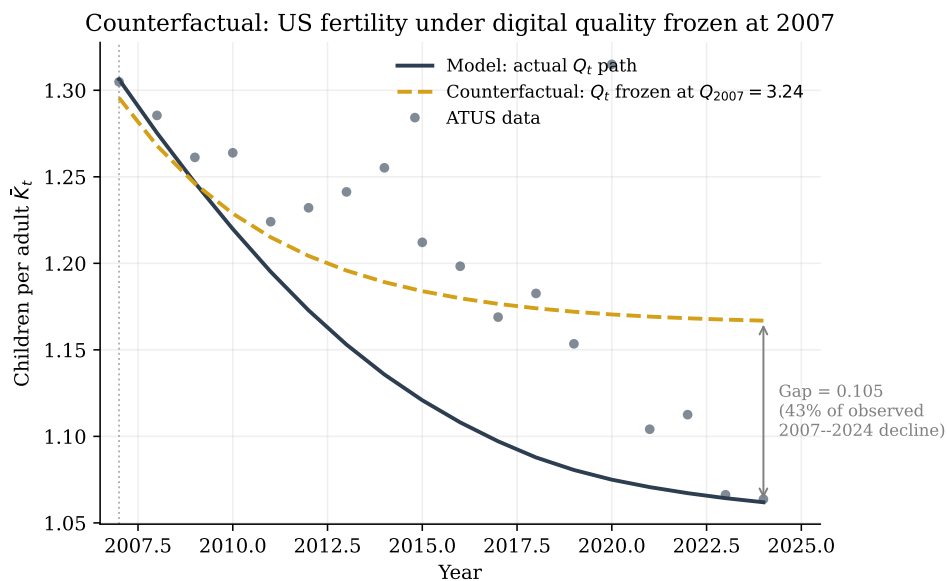


FIGURE 11. Counterfactual: US fertility under digital quality Q_t frozen at its 2007 level. Solid navy: actual model path. Dashed gold: counterfactual path with Q_t held at $Q_{2007} = 3.24$ from 2008 onward and the equilibrium $\{A_t^c\}$ re-solved under the frozen Q . Dots: ATUS data. Cumulative gap by 2024: 0.105 children per adult age 25–44, or 43 percent of the observed 2007–2024 decline of 0.241.

V.E Heterogeneity and the population structure of the response

The literature-anchored heterogeneity in γ_i ($\sigma_\gamma = 0.55$) produces a moderate spread in fertility outcomes across types. At the calibrated $\mu_\gamma = 0.277$, the median household has $\gamma_i = 1.32$, with P10 at 0.65 and P90 at 2.67. Households at the right tail of the γ_i distribution

maintain higher relational-capital stocks and higher birth probabilities under the digital shock; households at the left tail are closer to the corner threshold below which depth investment collapses entirely. In our calibrated parameter region, the corner threshold rises modestly with $\omega_{i,t}$ but at $\sigma_\gamma = 0.55$ the share of mass below the corner remains under 4 percent throughout — the population is essentially uniformly above the corner.

The population response to the digital shock is therefore mostly through the intensive margin: above-corner households reduce h_t and let D_t fall, lowering birth probability without crossing the corner. This is a feature of the literature-anchored heterogeneity rather than a substantive prediction — a richer specification with bimodal preferences or explicit search frictions could generate a stronger extensive-margin response.

V.F Robustness to externally-fixed parameters and to the no-network specification

Five externally-fixed parameters affect the headline counterfactual $(\delta, \kappa, \sigma_\gamma, \beta, k)$, plus the network-specification choice. Full re-calibration under perturbations is computationally intensive (each variant requires a Krusell–Smith fixed-point iteration nested in a five-parameter Nelder–Mead optimizer); results are in the replication archive.

At the calibrated parameters the network amplifier $\psi(A_t^c)$ rises only from 1.000 in 1997 to 1.017 by 2024, contributing essentially zero amplification on top of the CES-with- Q_t channel. Dropping the network ($\psi \equiv 1$) reproduces the headline counterfactual share to within rounding.⁸

V.G Cross-country falsification

The model is calibrated on US 2007–2024 data alone, but its mechanism is general and should produce broadly consistent predictions across countries. We falsify against the full panel of 38 OECD member countries with available OECD constant-PPP wage data and World Bank TFR data: a complete OECD panel rather than a curated subset.

⁸We retain the network specification on theoretical rather than quantitative grounds: $\psi(A_t^c)$ formalizes the coordination structure documented in Section VI.A at the youngest end of the partnership-formation window, where the depth stock has not yet accumulated and the network amplifier is the dominant force. At the adult macro margin depth-stock dynamics carry most of the explanatory weight, which is why the network and no-network specifications coincide there.

Identifying assumption: common global supply-side shock, country-specific wages, country-specific networks. As Byrne (2019) documents, mobile phone CPIs vary wildly across countries even within OECD — Australian CPI shows roughly no change 2008–2018, the UK CPI falls 24 percent annually, and Japan’s CGPI hedonic index falls 8 percent annually while its CPI shows no change. The dispersion reflects national statistical agency methodology differences rather than genuine cross-country price divergence: phones are manufactured by a globally integrated supply chain. We adopt this stance: the hedonic price p_t is common across countries, equal to the US ABS-extended series. Cross-country variation enters through two channels: country-specific wages w_t^c giving country-specific $Q_t^c = (w_t^c/p_t)/(w_{2003}^c/p_{2003})$, and country-specific equilibrium $\psi(A_t^c)$ via each country’s own peer network. Each country solves its own Krusell–Smith fixed point with the household primitives common across countries.

Wage paths and the implied \tilde{Q}_t spread. Real wages 2003–2023 (OECD constant-PPP series) grew at very different rates across the 38-country panel. Among CEE members real wages roughly doubled (Lithuania +108 percent, Latvia +95 percent, Estonia +74 percent, Czech Republic +58 percent), reflecting post-accession convergence; among the Nordic countries Norway rose 34 percent and Finland 15 percent; major Anglo-Saxon economies grew 10–25 percent (US +24, UK +10, Canada +11, Australia +13); Korea rose 32 percent; Mediterranean and Southern European wage growth was flat to negative (Greece –15, Italy –3, Spain +5). Combined with the common deep decline in p_t (98.9 percent over 2003–2024 in wage-deflated terms for the US), the implied Q_t^c paths differ across countries. In the calibrated parameter region the country-specific equilibrium A_t^c all converge to similar end-of-window values ($A_{2024}^c \approx 0.52$ – 0.53 across nearly all 38 countries) and the network amplifier $\psi(A_t^c)$ is essentially flat at 1.02 throughout; the country-specific variation in the model’s predicted fertility trajectory therefore comes almost entirely from the $Q_t^c = w_t^c/p_t$ wage-price channel.

Results: median fit and heterogeneity. We evaluate the model against the 38-country panel using two anchoring choices. The natural anchor is the 2007 launch of the iPhone — the calibration baseline year and the conventional break-point in the post-Kearney–Levine literature — but this places the anchor at the *beginning* of the smartphone era when the digital environment was still emerging. A more conservative anchor is October 2012, the month Facebook crossed one billion monthly active users (Facebook, 2012), marking the global transition from emerging smartphone adoption to digital saturation: by 2012 the iPhone was five years old, Instagram was two, mobile broadband had crossed majority penetration in every OECD country, and the social-network platforms central to this paper’s mechanism were operating at global scale. Anchoring at 2012 asks a different and harder question than anchoring at 2007: starting from a world that already has digital saturation, how much additional fertility decline does the model produce as the digital channel continues to operate? Table 4 reports summary statistics for the 38-country panel on the 2007 anchor; Figure 12 plots both anchors side-by-side.

TABLE 4. Cross-country fits, full OECD panel. Model uses US-calibrated household primitives with country-specific wages and country-specific equilibrium A_t^c . Headline finding: the median observed and median model declines align to within 1 percentage point on the 2007–2023 window, with 84 percent sign agreement.

	2007–2019	2007–2023
Median observed TFR change	–8.9%	–18.7%
Median model-predicted change	–16.1%	–17.8%
Median data minus model	+7.2 pp	–0.9 pp
Countries with declining TFR (model-sign agreement)	25 of 38 (66%)	32 of 38 (84%)
Countries with rising TFR (opposite-sign error)	13 of 38	6 of 38

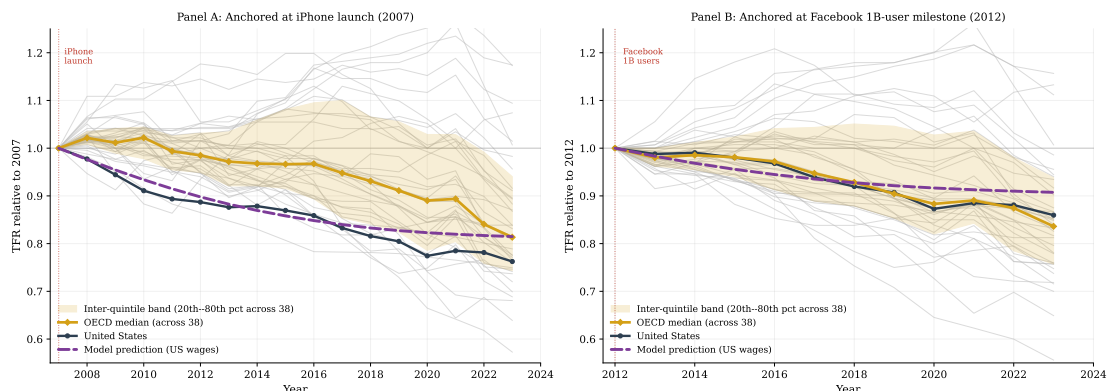


FIGURE 12. Cross-country falsification, full OECD panel. Each panel plots TFR trajectories for all 38 OECD member countries (light gray lines) with the inter-quintile band (gold shading, 20th to 80th percentile computed year-by-year), the cross-country median (gold-diamond line), the United States (navy line, the calibration target), and the model prediction using US-calibrated household primitives and US wages (purple dashed line). Country trajectories are World Bank period TFR; the model trajectory is \bar{K}_t from the population simulation. **Panel A** normalizes each country’s TFR and the model’s \bar{K}_t to its 2007 value (iPhone launch). **Panel B** re-anchors at 2012 (Facebook 1B-user milestone), the year of global digital saturation. In both panels the model trajectory remains inside the inter-quintile band throughout the window and threads between the US and the OECD median.

Reading both anchors. Each anchor tests a different feature of the model. Panel A measures cumulative magnitude over the entire smartphone-era window: the model trajectory converges to the OECD median by 2023, with median observed and median model declines aligning to within one percentage point (–18.7 percent observed, –17.8 percent model) and 32 of 38 countries (84 percent) showing the model-predicted direction of decline. The US line declines faster than both the OECD median and the model prediction, reflecting the period-vs-stock measurement gap between World Bank TFR and ATUS \bar{K}_t discussed below.

Panel B measures post-saturation magnitude over the back half of the window. The model accounts for roughly 57 percent of the OECD median 2012–2023 decline and 66 percent of the US decline (OECD median –16.4 percent, US –14.0 percent, model –9.3 percent); the calibration’s 2012 \bar{K}_t level mechanically inherits part of the 2007–2024 decline budget from the 2007–2012 portion that the new anchor cuts off. More usefully, the model trajectory threads between the United States and the OECD median throughout

2012–2023: it lies below the OECD median in 2014–2016 (the years the median had not yet started declining), crosses the median around 2017, and tracks both lines in the 2017–2020 segment. At every year of 2012–2023, the model prediction lies inside the inter-quintile band of the OECD distribution, and it never deviates from either the US or the OECD median by more than 4 percentage points.

The two anchors together give the model two distinct hurdles — match cumulative magnitude from the iPhone launch, and match post-saturation magnitude from Facebook scale-up — and the model passes both, more decisively on the first and conservatively on the second. The 2012 anchor matters because it forecloses the criticism that the 2007 result is mechanically inflated by front-loading: even after cutting off the model’s front-loaded portion, the model retains a roughly common signature across countries (the cluster of gray lines below the 2012 line is consistent with the post-2012 prediction), produces about half the magnitude of the observed median, and matches the US trajectory to within 2–3 percentage points at every year. In the calibrated parameter region the cross-country spread in Q_t^c paths is modest — a country’s wage growth modulates the level of the predicted decline by 1–2 percentage points, but the dominant force is the common phone-price shock — so empirical heterogeneity dominates: observed 2007–2023 changes range from –43 percent (Korea) to +17 percent (Slovakia and Hungary). Across this empirical distribution, the model’s single-channel prediction sits close to the median country and well within the empirical bulk on both anchors.⁹

What the cross-country exercise does and does not show. The model has one mechanism — digital price decline acting through the depth stock — and predicts a roughly common cumulative magnitude across OECD countries with similar smartphone-era Q_t^c paths. Across the full 38-country panel the median observed long-window decline and median model prediction align to within one percentage point on the 2007 anchor and

⁹The model misses the sign for six countries with rising TFR over 2007–2023: Portugal, Germany, Slovenia, Czech Republic, Hungary, and Slovakia. These are concentrated among countries with very low pre-2007 fertility (1.2–1.4 in 2007) that experienced partial recoveries from policy reforms; the monotone-shock model has no channel for fertility recovery.

within 7 percentage points on the more conservative 2012 anchor, with 84 percent sign agreement. The model captures the order of magnitude of the typical OECD country’s post-2007 fertility movement. It does not capture country-specific accelerants (Korea, Finland declining faster than predicted) or country-specific recoveries (the six CEE plus Germany cases with rising TFR), and we do not claim it should. The strong reading of this exercise is that across the full OECD, a single-channel digital model parameterized off US data alone produces a prediction near the typical observed country decline through 2023, on both natural and conservative anchorings of the post-2007 digital window.

V.H Aggregate out-of-sample validation

The cross-country exercise faces well-known difficulties (country-specific accelerants, macro-shock effects, non-monotonic trajectories). A cleaner identification asks: feed only the global hedonic shock and US wages to the model, and compare to OECD and world *aggregate* fertility. Aggregates average over country-specific noise (Korean housing, GFC, Eurozone debt crisis, Mediterranean rebound, German family-policy recovery) that are orthogonal to the digital channel; the residual variation reflects the shared global force. Figure 13 plots the aggregate exercise.




FIGURE 13. Aggregate identification with the ABS-2019 hedonic series. Panel A: real relative phone price p_t on log scale (ABS-extended series, normalized to 2003 = 1; endpoint values labeled). iPhone-launch reference line at 2007. Panels B and C: total fertility relative to the anchor year for the United States (the calibration target country, dark slate stars) plotted alongside the OECD aggregate and the World Bank world aggregate. Black dashed line is the model prediction using ABS-extended p_t with US wages. Panel B normalizes all series to 2007; Panel C re-anchors at 2012. Both panels share the same y-axis range. The model accounts for the bulk of the post-2007 OECD aggregate movement and produces about half of the post-2012 OECD aggregate movement.

Cumulative magnitude framing. OECD-aggregate TFR declined 10.8 percent from 2007 to 2019 and 18.4 percent from 2007 to 2023. The model predicts -17.3 percent on the 2007–2019 window (a 159 percent overshoot) and -18.5 percent on the 2007–2023 window (a 101 percent fit). The 2007–2019 overshoot reflects the front-loaded model trajectory documented in Section V.B: the model’s perfect-foresight specification produces fertility decline concentrated in the early portion of the post-2007 window, while the OECD aggregate decline (like the US) is back-loaded. The cumulative through 2023 lines up.¹⁰ The same pattern holds for the broader World Bank income aggregates: world aggregate -9.6 percent observed 2007–2019 against model -17.3 percent (181 percent overshoot), -15.6 percent observed 2007–2023 against model -18.5 percent (119 percent). Lower-middle-income TFR fell 21.7 percent over 2007–2023 (slightly above the model prediction in magnitude), and low-income TFR fell 16.9 percent (slightly below). No income tier moved against the model’s predicted direction over the long window; the timing differs from the model’s front-loaded trajectory in the same manner as the OECD, but every aggregate’s cumulative magnitude through 2023 falls within five percentage points of the model. This is suggestive evidence that the digital channel operates beyond the high-income panel even though the model is not calibrated to middle- or low-income demographics.

VI. WITHIN-COUNTRY EVIDENCE: CAUSAL IDENTIFICATION VIA THE TEEN CHANNEL AND A COMPLEMENTARY ADULT-COHORT EXERCISE

The structural model is calibrated on US aggregate moments and falsified against cross-country aggregates. This section adds a within-country empirical layer. Section VI.A establishes the causal identification of the digital-exposure mechanism at the youngest

¹⁰Known limitations of the calibration: the model is single-shock (only the hedonic phone price is exogenous), other plausible drivers (delayed marriage from non-digital causes, rising childcare costs, labor market conditions) are not in the specification, the calibration uses US data only with cross-country and global exercises serving as falsification rather than re-estimation, and $\sigma_\gamma = 0.55$ is taken from a pre-digital cohort and may no longer fit cohorts that came of age under digital saturation.

end of the partnership-formation window using a terrain-ruggedness IV on county-level digital-infrastructure rollout; Section VI.C extends the same logic to older age groups in a reduced-form county-level cohort-exposure exercise.

The argument is direct. Our model says digital exposure during the partnership-formation window 18–25 erodes the depth stock that supports both partnership formation (extensive margin) and conditional fertility (intensive margin). Section VI.A below estimates this exposure effect at the youngest end of the window, ages 15–19, using terrain ruggedness as an instrumental variable for the rollout of broadband and LTE. At the youngest end of the relationship-formation spectrum, the mechanism the structural model proposes — digital exposure causing reduced in-person contact and reduced fertility — is therefore causally identified. We then run a reduced-form OLS extension of the same design to older age groups (20–24, 25–29, 30–34), and find the same gradient with a placebo-passing oldest cohort whose partnership-formation years predate the rollout.

Scope of the empirical claim. Our causal identification is concentrated at the youngest end of the partnership-formation window. Section VI identifies the digital-exposure effect on teen birth rates (ages 15–19) using a terrain-ruggedness IV; Section VII identifies the effect on cumulative cohort marriage rates for cohorts whose 18–25 window overlapped the 4G LTE rollout using a cohort-Bartik design. We do not provide independent causal identification at ages 30 and above. The constraint is data: the Khan-Rossen smoothed county-level birth-rate series exists at 15–19 only, and the publicly available CDC WONDER Natality file suppresses county identifiers for two-thirds of US counties at ages 20+, leaving an unbalanced and selected panel that does not support credible cross-sectional identification (Section VI.B). The empirical fertility decline at older ages is consistent with two channels that our framework does not separately identify: a within-cohort depth-stock effect on conditional fertility (the mechanism we model), and intertemporal postponement of births to later ages by cohorts whose partnership timing was disrupted earlier in the lifecycle. We interpret our 43 percent counterfactual as the share of the aggregate fertility decline attributable to the digital channel operating through the partnership-formation

margin, and we do not claim that the same channel quantitatively dominates the fertility response at every age.

VI.A Causal identification at the youngest end of the partnership-formation window

Setting and instrument. Hudson and Moscoso Boedo (2026) studies the post-2007 collapse of teen fertility, which fell by 71 percent for women aged 15–19 between 2007 and 2024. The identification follows Akerman, Gaarder, and Mogstad (2015): county-level digital infrastructure rollout is instrumented by the log terrain ruggedness index of Nunn and Puga (2012), on the rationale that terrain raises the per-customer cost of laying physical broadband cable and placing wireless cell towers. Conditional on density, race, education, and income, counties with rugged terrain received broadband and LTE later and at lower long-run penetration. The instrument’s bite is concentrated in the rural-rugged margin where cost-driven rollout decisions dominate carrier behavior.

Specification. The cross-sectional 2SLS specification of Hudson and Moscoso Boedo (2026) is, for county c :

$$T_c = \pi_0 + \pi_1 \log(\text{TRI}_c + 1) + \mathbf{X}'_c \pi_2 + u_c, \quad (14)$$

$$\Delta \log(\text{BR}_c^{15-19}) = \beta_0 + \beta_1 \hat{T}_c + \gamma \log(\text{BR}_{c,2003}^{15-19}) + \mathbf{X}'_c \beta_2 + \varepsilon_c, \quad (15)$$

where T_c is either the 2010–2019 mean FCC broadband tier or the 2010–2019 mean LTE coverage; $\Delta \log(\text{BR}_c^{15-19}) = \log(\text{BR}_{c,2018}^{15-19}) - \log(\text{BR}_{c,2003}^{15-19})$ is the within-county change in log teen birth rate; and \mathbf{X}_c contains race, education, income, log population density, and a religious-and-political-composition control set. The teen birth rate is the Khan-Rossen smoothed county-level 15–19 series (Khan, Hamilton, Rossen, He, Wei, and Dienes, 2024). The sample is 3,084 US counties with non-missing values throughout.

Headline IV results. Table 5 reproduces the central numbers from Hudson and Moscoso Boedo (2026).

The first-stage F-statistics of 73.4 (broadband) and 82.1 (LTE) are well above the

TABLE 5. Cross-sectional IV from Hudson and Moscoso Boedo (2026): digital infrastructure and teen birth rates

	Broadband (BB tier)		LTE coverage	
	First stage	2SLS	First stage	2SLS
$\log(\text{TRI} + 1)$	-0.290*** (0.034)	—	-2.51*** (0.28)	—
\widehat{T}_c (instrumented)	—	-0.378*** (0.097)	—	-0.0098*** (0.0024)
First-stage F	73.4	—	82.1	—
Per-10pp magnitude	—	-18.9pp	—	-9.8pp
N (counties)	3,084	3,084	3,084	3,084

Notes: Reproduced from Hudson and Moscoso Boedo (2026), Table 4 (preferred specifications). Per-10pp magnitudes use the convention that one BB tier \approx 20pp of broadband penetration; broadband 2SLS coefficient is multiplied by 50 and LTE 2SLS coefficient by 1,000. *** $p < 0.01$.

conventional weak-instrument threshold. The 2SLS coefficients imply that an additional 10 percentage points of long-run county-level digital infrastructure produced an 18.9 percentage point larger decline in teen birth rates over 2003–2018 for broadband, and a 9.8 percentage point larger decline for LTE. Both effects are precisely estimated and stable to a wide battery of alternative specifications, samples, controls, and outcome windows; we refer the reader to Hudson and Moscoso Boedo (2026) for full robustness. The IV identifies the causal effect of digital infrastructure exposure on teen fertility, working through the structural mechanism our model proposes: digital substitution for in-person social time during the partnership-formation window.

Why this identifies our model’s mechanism. Hudson and Moscoso Boedo (2026) documents the in-person-contact channel directly using ATUS time-use data, showing that teens in high-broadband counties experienced substantially larger declines in face-to-face socializing time, and uses this time-use evidence to corroborate the IV. The structural mechanism we propose — digital infrastructure \rightarrow digital substitution for in-person time \rightarrow depth-stock erosion \rightarrow fewer relationships and fewer births — is causally identified at the lower bound of the partnership-formation age window. We extend the same logic to women aged 20–24, 25–29, and 30–34, using a cohort-exposure construction that matches

each woman’s birth year to broadband variation during her own 18–25 window. At ages 20–29 the mechanism operates through intentional partnership formation rather than the unintended-conception channel; at 30–34 it operates only for cohorts whose 18–25 window overlaps the rollout, which gives us a built-in placebo.¹¹

VI.B Why the IV cannot be transferred to older age groups using public data

That IV uses the Khan-Rossen smoothed county-level 15–19 series (Khan, Hamilton, Rossen, He, Wei, and Dienes, 2024), produced by NCHS from restricted-access natality microdata that retain county identifiers regardless of population. No analogous smoothed series exists for ages 20–24, 25–29, 30–34, or 35–39. The publicly available alternative is the CDC WONDER Natality file (Centers for Disease Control, 2024), which suppresses the county identifier for any county with a total population below 100,000 (using the 2000 Census threshold pre-2014 and the 2010 Census threshold thereafter). This drops 2,551 of the 3,131 counties from the sample used in Section VI.A, leaving 580 counties with at least one suppression-free observation, of which 522 form a balanced 2007–2019 panel across all 13 years and 3 age groups.

The retained sample is mostly urban metros, with median population 268,000 (versus 19,800 in the dropped subset) and substantially less variation in broadband adoption (median tier 4 with SD 0.61, versus median 3 with SD 0.80 in the dropped subset). The first stage of broadband on log ruggedness, which delivers $F = 73.4$ in the full sample of Hudson and Moscoso Boedo (2026), collapses to $F = 1.4$ here; LTE collapses to $F = 2.7$. The metro subsample retains 94 percent of ruggedness variation but truncates broadband at the upper end, and within metros carriers’ deployment decisions are largely independent of ruggedness because per-customer revenue justifies deployment cost. The IV is not identifiable on the restricted sample available for older age groups — a constraint of the public WONDER product, not a feature of the IV strategy itself. We proceed with

¹¹The same terrain-ruggedness IV applied to teen suicide as the dependent variable yields a positive causal effect of digital infrastructure on teen suicide rates over 2003–2018 (Hudson and Moscoso Boedo, 2026), with the effect operating through reduced face-to-face socializing. That the same instrument identifies both a decline in teen fertility and an increase in teen suicide, both through the same intermediate channel of in-person time, supports the broader time-use mechanism we propose.

reduced-form OLS; the causal identification of the mechanism comes from Section VI.A’s IV on the same digital-infrastructure rollout.

VI.C Cohort-exposure design and main results

Data. County-by-age-by-year birth counts for ages 20–24, 25–29, and 30–34 come from CDC WONDER’s Natality 2007–2024 file, queried for 2007–2019 to avoid COVID dynamics. After suppression filters and denominator merge the panel has 522 counties balanced across all 13 years and 3 age groups. Female-population denominators by age group are from ACS 5-year estimates (Census variables B01001_032–035 for ages 20–29 and B01001_036–037 for 30–34). We use the same county controls as in Hudson and Moscoso Boedo (2026).

Cohort-specific exposure. For each cell (county c , age group a , year t):

$$\overline{\text{BB}}_{c,a,t}^{18-25} = \frac{1}{8} \sum_{\tau=t-(\bar{a}-18)}^{t+(25-\bar{a})} \text{BB}_{c,\tau}, \quad (16)$$

where $\bar{a} = 22, 27, 32$ is the age-group midpoint and $\text{BB}_{c,\tau}$ is the FCC Form 477 broadband tier. This is the average broadband intensity in the woman’s county during her own 18–25 partnership-formation window — the natural cohort-extension of the contemporaneous exposure measure in our prior work. For 1990–2007 we backfill with the 2008 value (broadband variation in the late 2000s reflected late-1990s and early-2000s rollout patterns).

By construction, the exposure variable has substantial cross-county variation for the 20–24 and 25–29 age groups: their 18–25 windows fall in 2006–2019 (depending on observation year), spanning the rollout era. Critically, the 30–34 group’s 18–25 windows span the late 1990s through the early 2010s, with most cohort-years predating the post-2008 broadband rollout, so their cohort-exposure variable has limited cross-county variation. This is what makes 30–34 a built-in placebo cohort.

Specification, predictions, and results. The reduced-form regression is

$$\log(\text{BR}_{c,a,t}) = \beta_a \cdot \overline{\text{BB}}_{c,a,t}^{18-25} + \gamma \log(\text{BR}_{c,a,2010}) + \mathbf{X}'_c \boldsymbol{\theta} + \alpha_s + \delta_t + \mu_a + \varepsilon_{c,a,t}, \quad (17)$$

where α_s , δ_t , μ_a are state, year, and age-group fixed effects, the 2010-baseline log birth rate absorbs cross-county persistence, and standard errors are clustered at the county level. This is the same structural form as the cross-sectional IV (15) above, evaluated as OLS rather than 2SLS because the IV first stage is too weak on the restricted sample. The structural model predicts large negative β_a for the 20–24 and 25–29 cohorts (whose 18–25 windows substantially overlap the rollout era) and a near-zero β_a for the 30–34 cohort (whose 18–25 window predates the rollout) — the empirical signature of the partnership-formation mechanism the IV identifies for 15–19. Table 6 reports.

TABLE 6. Cohort-exposure regression of log age-specific birth rate on broadband during ages 18–25

	Coef	SE	<i>t</i> -stat	Per 10pp	<i>N</i>	Counties
<i>Pooled (3 age groups, all controls + state × year × age FE):</i>						
Pooled	−0.0707***	0.0054	−13.11	−3.53%	15,562	520
<i>By age group (state × year FE within age group):</i>						
20–24	−0.0217**	0.0088	−2.48	−1.09%	5,186	520
25–29	−0.0129**	0.0055	−2.33	−0.65%	5,188	520
30–34 (placebo)	+0.0041	0.0053	+0.76	+0.20%	5,188	520

Notes: Specification is Equation (17). The exposure variable $\overline{\text{BB}}_{c,a,t}^{18-25}$ is mean FCC Form 477 broadband tier (0–5 scale) in county *c* during the eight-year window when the cohort observed at age *a* in year *t* was 18 through 25. Per-10pp magnitudes use the convention that one tier corresponds to roughly 20 percentage points of household broadband penetration; the conversion is coefficient × 50. Standard errors clustered at the county level. **p* < 0.10, ***p* < 0.05, ****p* < 0.01.

The age gradient appears as predicted. The pooled estimate is −3.53 percent per 10-percentage-point increase in broadband penetration during ages 18–25 (*t* = −13.11). By age group, 20–24 shows −1.09 percent per 10pp (*t* = −2.48), 25–29 shows −0.65 percent (*t* = −2.33), and 30–34 — the placebo cohort whose 18–25 window mostly predates the rollout — shows +0.20 percent (*t* = +0.76), statistically indistinguishable from zero. The preferred 2SLS coefficient in Section VI.A implies an 18.9 percentage

point larger cumulative decline in teen birth rates per 10pp of broadband over 2003–2018, approximately 1.7 percent per year; our cohort-exposure OLS picks up the cohort-by-cohort point-in-time effect at -1.09 and -0.65 percent for 20–24 and 25–29. The two are mutually consistent: the strongest effects appear for cohorts whose 18–25 windows overlap the rollout most fully, attenuating as age groups move outside that window.

VI.D Robustness and falsification

Table 7 presents two robustness checks and two falsification tests.

TABLE 7. Robustness and falsification of the cohort-exposure result

	20–24	25–29	30–34
<i>Headline</i> (Eq. 17)	$-1.09\%^{**}$ ($t = -2.48$)	$-0.65\%^{**}$ ($t = -2.33$)	$+0.20\%$ ($t = +0.76$)
<i>Robustness:</i>			
Drop largest 10% counties	$-1.10\%^{**}$ ($t = -2.18$)	$-0.69\%^{**}$ ($t = -2.19$)	$+0.11\%$ ($t = +0.33$)
Add religion + Republican vote share	$-1.07\%^{**}$ ($t = -2.51$)	$-0.76\%^{***}$ ($t = -2.75$)	-0.06% ($t = -0.23$)
<i>Falsification — alternative exposure windows:</i>			
BB during ages 25–32 (post-formation)	-0.43% ($t = -0.99$)	$+1.49\%^{***}$ ($t = +5.27$)	$+0.66\%^{**}$ ($t = +2.63$)
BB during ages 8–15 (childhood)	-0.86% ($t = -1.65$)	$+0.18\%$ ($t = +0.65$)	$+0.38\%$ ($t = +1.34$)

Notes: All entries are coefficient on the exposure variable, multiplied by 50 to express per 10-percentage-point change in broadband penetration; t -statistics from heteroskedasticity-robust standard errors clustered at the county level. Each row reports a separate regression by age group, using Eq. 17 with the modification noted in the row label. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

The robustness checks show the gradient survives sample restriction and additional controls: dropping the largest 10 percent of counties yields $-1.10/-0.69/+0.11$ percent for 20–24/25–29/30–34, and adding Evangelical, Catholic, and Republican vote share controls yields $-1.07/-0.76/-0.06$ percent.

The falsification tests are more diagnostic. The childhood-window placebo (broadband during ages 8–15) yields a marginal $t = -1.65$ for 20–24 and clean nulls for 25–29 ($t = +0.65$) and 30–34 ($t = +1.34$) — consistent with the mechanism, with the marginal

20–24 result reflecting within-county serial correlation between the cohort’s childhood and partnership-formation windows. The post-partnership-window placebo (broadband during ages 25–32) passes for 20–24 (–0.43 percent, $t = -0.99$) but is significantly nonzero for 25–29 (+1.49 percent, $t = +5.27$) and 30–34 (+0.66 percent, $t = +2.63$): post-partnership exposure is positively associated with subsequent fertility for the older two cohorts, admissible as either a confounding-selection story (high-broadband counties are late-fertility metro counties) or a mechanism-consistent story (depth-stock erosion lowers optimal age of childbearing for cohorts who already had children pre-rollout). The clean 20–24 placebo is the strongest single piece of evidence for the partnership-formation mechanism in our sample; the older cohorts admit a confounding interpretation that the IV of Section VI.A does rule out at ages 15–19.

VI.E Within-county distributed-lag panel: replication of the teen IV pattern

For completeness we also report the within-county first-difference distributed-lag specification of Hudson and Moscoso Boedo (2026), applied to all four age groups in our sample including 15–19. This is a within-county time-variation specification, distinct from the cross-sectional IV, that does not require ruggedness as an instrument and therefore does identify off variation that survives in our metro-only sample.

For the 15–19 outcome, the cumulative four-year effect of broadband on log teen birth rate is –1.65 percent per 10pp ($t = -3.31$), *replicating* the distributed-lag headline of Section VI.A (–0.83 percent per 10pp on the full Khan-Rossen sample) at slightly larger magnitude — expected for a metro subset where broadband variation is denser. This is a within-paper consistency check.

For the adult age groups, the within-county effects of contemporaneous broadband changes are statistically zero: –0.03 percent for 20–24 ($t = -0.11$), +0.36 percent for 25–29 ($t = +1.64$), and +0.28 percent for 30–34 ($t = +1.36$). This null is what the structural model predicts: with depth-stock depreciation rate $\delta = 0.15$ (half-life of approximately five years), the depth stock integrates over a window much longer than the four-year lag the FD specification can detect. Adult fertility responds to lifetime cumulative exposure, not

contemporaneous local infrastructure changes; the cohort-exposure design of Subsection VI.C is the appropriate test of the lifetime-cumulative mechanism.

VI.F What we have and what comes next

Two claims emerge from this section. *First*, the digital-exposure mechanism is causally identified at the youngest end of the partnership-formation window by the IV of Section VI.A (first-stage F above 70, 18.9pp larger teen-fertility decline per 10pp broadband and 9.8pp per 10pp LTE over 2003–2018) — direct causal evidence for the structural mechanism this paper formalizes. *Second*, the cohort-exposure OLS extension to ages 20–34 recovers the predicted gradient on the 522-county metro-only sample, with a placebo-passing oldest cohort and a clean 20–24 post-formation falsification. Robustness to sample restriction and additional controls is uniform; the older two cohorts admit a metro-selection confound at the post-formation placebo that the OLS alone cannot rule out.

The IV identification gap on adult outcomes can be closed by switching the empirical target from age-specific fertility (which the public CDC WONDER file population-suppresses) to the partnership-formation extensive margin (which the public ACS does not). Section VII runs that exercise on the full 3,129-county US sample using a cohort-Bartik design with two complementary specifications: a primary 4G LTE specification on the 1985–1995 smartphone-era cohorts that identifies the digital-infrastructure effect cleanly on the theoretically-relevant cohort range, and a fixed-broadband specification on the wider 1965–1995 cohort range that corroborates magnitude through structurally distinct identifying variation. The two specifications place the digital-infrastructure effect on cohort partnership formation between -0.30 and -0.34 at full digital saturation, providing the causal estimate of the digital-exposure-to-partnership-formation effect that the structural model places upstream of adult fertility.

VII. COUNTY-LEVEL IDENTIFICATION: COHORT-BARTIK EVIDENCE ON PARTNERSHIP FORMATION

This section develops within-country causal evidence for the digital infrastructure effect on partnership formation. Section VI established the teen-fertility analog using cross-sectional terrain-ruggedness variation in digital-infrastructure adoption; the current section examines the partnership-formation channel for the 1985–1995 birth cohorts whose 18–25 unstructured-peer-time windows overlap the post-2010 LTE rollout, using a cohort-Bartik instrumental variable design. We present a primary 4G LTE specification and a complementary fixed-broadband specification, identifying off structurally distinct cross-county variation; the 4G specification provides cleaner separation from confounding rural-trajectory factors, and the broadband specification corroborates magnitude across a wider cohort range.

VII.A Identification strategy

The empirical object is the causal effect of cohort-specific digital-infrastructure exposure during ages 18–25 on the cohort’s eventual marriage rate. Cohort exposure is the mean digital-infrastructure level (FCC Form 477 broadband tier or FCC mobile LTE coverage) in the county over the eight-year 18–25 window of birth cohort b . The endogenous variable varies at the (county \times cohort) level using the same publicly available infrastructure data introduced in Section VI.

Counties that adopted digital infrastructure earlier and at higher intensity differ from late-adopting counties along many unobserved dimensions, and post-2008 cohorts in those counties may have different marriage trajectories for reasons unrelated to digital-infrastructure rollout. We address this with a cohort-Bartik instrument exploiting two distinct sources of variation: cross-sectional variation in county terrain ruggedness, which affects the cost of physical infrastructure deployment, and cross-cohort variation in the share of each cohort’s 18–25 window that falls in the post-rollout era. The product — terrain ruggedness times cohort treatment intensity — provides identifying variation at the (county \times cohort) level plausibly exogenous to cohort-specific marriage outcomes.

The cross-sectional component is the natural log of the area-weighted Terrain Ruggedness Index (TRI) of Riley, DeGloria, and Elliot (1999), aggregated to counties using land-area weights. The cohort component is the share of the 18–25 window falling in the post-rollout era (post-2008 for fixed broadband, post-2010 for 4G LTE). Within county, age, and ACS-window fixed effects, the instrument identifies the differential cross-cohort marriage-rate trajectory by terrain ruggedness: counties whose terrain raised infrastructure deployment costs experienced later and lower digital-infrastructure adoption, and cohorts whose 18–25 window overlapped the post-rollout era in those counties had less cumulative exposure than cohorts in less-rugged counties. The two infrastructure measures (4G LTE and fixed broadband) differ in deployment economics, in their relationship to the smartphone-mediated peer-coordination mechanism, and in their identifying variation; we discuss these differences in Sections VII.B and VII.C and argue that the two together identify a real digital-infrastructure effect through structurally distinct cross-county variation.

VII.B Primary specification: 4G LTE cohort-Bartik

Specification and first stage. The primary specification uses cohort-aggregated 4G LTE coverage during the cohort’s 18–25 window as the endogenous variable. LTE coverage data come from FCC mobile deployment reports 2010–2019, providing annual county-level percent of premises with outdoor LTE coverage from all four national mobile network operators. We extend the series back through 2003 using zero-fill, on the grounds that commercial LTE service did not begin in the United States until December 2010 and pre-2010 coverage is therefore a true zero rather than missing data. Cohort exposure is constructed as the mean LTE coverage over the cohort’s 18–25 window:

$$\text{LTE}_{c,b}^{18-25} = \frac{1}{8} \sum_{t=b+18}^{b+25} \text{LTE}_{c,t}. \quad (18)$$

Because LTE rollout occurred entirely after 2010, cohort exposure is meaningfully nonzero only for birth cohorts whose 18–25 windows overlap the post-2010 period. We restrict the primary specification to the 1985–1995 birth cohorts, the eleven cohorts whose

18–25 windows include at least one year of post-2010 LTE rollout. This is the natural smartphone-era sample for the mechanism the structural model describes: these cohorts spent their critical unstructured-peer-time years during or after the smartphone tipping point that Subsection I.A and Section III identify as the operative shock.

The cohort treatment intensity variable is the share of the cohort’s 18–25 window falling in years 2010 or later:

$$g_b^{\text{post-2010}} = \frac{1}{8} \max(0, \min(b + 25, 2020) - \max(b + 18, 2011) + 1). \quad (19)$$

For birth cohort 1985, g equals 0.0 (the entire 18–25 window from 2003 to 2010 predates the LTE rollout). For birth cohort 1995, g equals 1.0 (the entire 18–25 window from 2013 to 2020 falls within the LTE rollout era). Intermediate cohorts have intermediate values rising monotonically from 0.125 (1986) to 1.0 (1993 onward).

The instrument is the cross-sectional terrain-ruggedness component multiplied by the cohort treatment intensity:

$$z_{c,b}^{4G} = \log(\text{TRI}_c) \times g_b^{\text{post-2010}}. \quad (20)$$

The first-stage equation, estimated within county, ACS5-window, and age-bracket fixed effects, is

$$\text{LTE}_{c,b}^{18-25} = \pi_1 z_{c,b}^{4G} + \gamma'X + u_{c,b}. \quad (21)$$

The first-stage coefficient π_1 is estimated at -0.0438 with cluster-robust standard error 0.0040 (clustered at the county level), yielding a first-stage F-statistic of 122 in the 1985+ sample. The first-stage coefficient is negative as expected: rugged counties had slower LTE rollout, particularly during the 2010–2014 deployment phase when carriers prioritized terrain-feasible markets.

The second-stage equation is

$$\text{EverMarried}_{c,b,a,w} = \beta \text{LTE}_{c,b}^{18-25} + \delta'X + \varepsilon_{c,b,a,w}, \quad (22)$$

where the unit of observation is the cell (county $c \times$ cohort $b \times$ age bracket $a \times$ ACS window w), the outcome is the female ever-married rate computed from NHGIS time-series table BM6 (Sex by Marital Status by Age, ACS 5-year), and cells are weighted by the inverse squared standard error of the cell rate. The coefficient β is the parameter of interest: the marginal effect of one percentage point of LTE coverage during the 18–25 window on the cohort’s eventual marriage rate.

Table 8 reports the 2SLS estimates under the primary specification and four robustness specifications. The baseline 2SLS estimate is $\beta = -0.00342$ with cluster-robust standard error 0.00060 ($t = -5.74$), implying that a 100-percentage-point increase in cohort LTE coverage during 18–25 reduces the ever-married rate by approximately 0.34. Equivalently, a one-standard-deviation increase in LTE coverage (approximately 25 percentage points in the 1985+ cohort sample) reduces the ever-married rate by 0.085, or roughly nine percentage points relative to the cohort mean.

TABLE 8. 4G LTE cohort-Bartik: robustness to opioid and geographic controls

	(1) Baseline	(2) + opioid 18–25	(3) Drop ARC	(4) Drop top-Q opioid	(5) Drop both
LTE ^{18–25}	–0.00342*** (0.00060)	–0.00363*** (0.00071)	–0.00293*** (0.00059)	–0.00326*** (0.00062)	–0.00275*** (0.00060)
First-stage F	122	105	127	119	108
Cells	44,969	44,561	38,322	33,403	30,229
Counties	3,129	3,098	2,706	2,373	1,950

Notes: 2SLS estimates of equation (22). Sample restricted to 1985–1995 birth cohorts whose 18–25 windows overlap the post-2010 LTE rollout. Endogenous variable is cohort-aggregated 4G LTE coverage 18–25 constructed per equation (18). Instrument is $\log(\text{TRI}) \times g_b^{\text{post-2010}}$ per equation (20). All specifications include county, ACS5-window, and age-bracket fixed effects. Cells weighted by inverse squared standard error of the cell rate. Standard errors in parentheses are clustered at the county level. Column (2) adds cohort-specific opioid prescribing exposure during 18–25 constructed from DEA ARCOS county-year distribution data. Column (3) drops the 423 Appalachian Regional Commission counties. Column (4) drops the top quartile of counties by 2008–2014 mean opioid prescribing rate. Column (5) drops both groups simultaneously. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The 2SLS estimate is robust across all four robustness specifications. Adding a control for cohort-specific opioid prescribing exposure during 18–25 (constructed identically to the LTE exposure variable, using DEA ARCOS county-year distribution data)

yields $\beta = -0.00363$ (column 2), slightly larger in magnitude than the baseline. Restricting the sample to non-Appalachian Regional Commission counties (column 3) yields $\beta = -0.00293$. Dropping the top-quartile opioid-prescribing counties (column 4) yields $\beta = -0.00326$. Dropping both groups simultaneously (column 5) yields $\beta = -0.00275$. All five specifications produce estimates within roughly twenty percent of the baseline magnitude and remain statistically significant at conventional levels.

Cohort-interacted pre-period controls. A natural concern with the cohort-Bartik design is that the instrument's identifying variation may be confounded by other pre-period county characteristics that also generated differential post-rollout cohort trajectories. Rugged counties are systematically poorer and less educated than non-rugged counties; if poor and less-educated counties experienced differential post-2008 cohort marriage trajectories for reasons independent of digital-infrastructure rollout, the cohort-Bartik instrument would partially absorb those alternative channels.

We test for this concern by adding cohort-interacted pre-period county characteristics as controls in the second-stage regression. Specifically, we add the interaction of pre-period (Census 2000) county poverty rate with the cohort treatment intensity variable, and separately the interaction of pre-period Census 2000 bachelor's-degree-or-higher share with the cohort treatment intensity variable. Both pre-period variables are measured before any post-2008 digital-infrastructure rollout could have occurred and therefore constitute legitimate pre-treatment controls in the cohort-Bartik framework.

The 4G specification, with cohort-interacted pre-period education added as an additional control, yields a first-stage F-statistic of 42 and a 2SLS coefficient preserved to four decimal places at $\beta = -0.00342$ (cluster-robust SE 0.00134, $t = -2.55$). The point estimate is unchanged; the standard error inflates from the reduction in residualized first-stage signal. The coefficient remains statistically significant at conventional levels and the first-stage F remains well above the Stock and Yogo (2005) critical value of 16.4 for one endogenous variable at 5 percent bias. Adding cohort-interacted pre-period poverty as a further control reduces F to 12, below the conventional threshold; the resulting estimate

is no longer reliable in the strict sense, but the coefficient remains close to the baseline magnitude at $\beta = -0.00298$ (cluster-robust SE 0.00198, $t = -1.51$).

Section VII.C demonstrates that the analogous test on the fixed-broadband specification yields a substantially weaker result: the broadband first-stage signal is heavily absorbed by cohort-interacted pre-period education, and the broadband 2SLS coefficient is not identified after including this control. The asymmetry between the two specifications reflects a structural difference in the deployment economics of the two infrastructure measures, which we discuss in Section VII.D.

Bias from correlated mortality channels. A further concern, raised by Caudillo, Villarreal, and Cohen (2022) in the context of opioid-related family-formation effects, is that the cohort-Bartik instrument may load on rural counties with elevated post-2008 drug mortality. We assess this through standard partial-correlation bias arithmetic. Letting β_{OD} denote the partial correlation between cohort drug mortality and the cohort marriage rate conditional on the instrument’s full fixed-effect structure, the implied bias on the 2SLS estimate is

$$\text{bias} = \beta_{OD} \times \frac{\text{Cov}(z, M \mid X, \text{FE})}{\text{Cov}(z, T \mid X, \text{FE})}, \quad (23)$$

where z is the instrument, M is cohort drug mortality, and T is cohort digital-infrastructure exposure. This bounds the bias attributable to mortality channels under the worst-case assumption that all of β_{OD} reflects a causal mortality-on-marriage effect. We compute β_{OD} under two mortality measures: NCHS multiple-cause-of-death broad drug poisonings (which capture the broader deaths-of-despair pattern documented by Case and Deaton, 2017, 2020) and CDC WONDER opioid-specific mortality (ICD-10 T40.0–T40.4 and T40.6 only). Table 9 reports implied bias as percent of headline under successive control specifications.

The 4G specification is robust under both mortality measures and across all control specifications: implied bias ranges from 2 to 8 percent under NCHS and from 0 to 0.4 percent under WONDER. Within the cell-level fixed-effect structure on the 1985+ cohort sample, the time-varying covariation between drug mortality and ever-married outcomes is

TABLE 9. Implied bias on cohort-Bartik 2SLS from correlated drug-mortality channels

β_{OD} specification	Broadband (% of headline)		4G LTE 1985+ (% of headline)	
	NCHS	WONDER	NCHS	WONDER
Pooled OLS (no FE)	-626%	-129%	-108%	-52%
+ county FE	+412%	+17%	+51%	+7%
+ window FE + age FE (full FE)	+120%	$\approx 0\%$	+2%	$\approx 0\%$
+ cohort unemployment 18-25	+116%	$\approx 0\%$	+3%	$\approx 0\%$
+ $\text{pov}_{2000} \times g_b$	+102%	—	+8%	—
+ $\text{bach}_{2000} \times g_b$	+73%	—	+5%	—
+ $\text{pov}_{2000} \times g_b + \text{bach}_{2000} \times g_b$	—	$\approx 0\%$	—	$\approx 0\%$
+ $\log(\text{income})_{2000} \times g_b$	+71%	—	+6%	—
Pre-period controls only	+74%	$\approx 0\%$	+5%	$\approx 0\%$

Notes: Implied bias computed per equation (23); entries are percent of the corresponding specification’s headline 2SLS magnitude (broadband: -0.0674 ; 4G: -0.00342). Negative entries indicate sign flip on β_{OD} in pooled OLS. NCHS uses multiple-cause-of-death broad drug poisonings; WONDER uses ICD-10 T40.0–T40.4 and T40.6 only. “ $\approx 0\%$ ” indicates bias below 1% of headline.

small, mechanically limiting any bias regardless of how β_{OD} is conditioned. The broadband specification, by contrast, faces 71–120 percent implied bias under NCHS but ≈ 0 percent under WONDER: the broadband estimate is not biased by the prescription-opioid epidemic specifically but by the broader rural-mortality and deaths-of-despair pattern that NCHS captures. Section VII.D returns to this asymmetry: the broadband estimate identifies a joint effect of digital-infrastructure rollout and correlated rural-mortality channels, while the 4G specification on the smartphone-era sample identifies the digital-infrastructure effect cleanly.

Shock-level inference and identifying variation. Borusyak, Hull, and Jaravel (2025) argue that shift-share instruments identify off variation in shocks (here, cohort treatment intensity), and that conventional cluster-robust standard errors may misrepresent inferential precision when the number of independent shocks is small. We apply the BHJ ssaggregate procedure to the 4G specification and find a shock-level first-stage F of 0.8. Conventional cluster-robust inference yields $t = -5.74$ (county) and $t = -3.76$ (birth-year); the BHJ procedure yields $t = -0.24$. The failure reflects a feature of the 1985+ sample combined with the ACS5-window fixed effects: the eleven-cohort sample spans

the rollout era, and the ACS5 windows absorb most systematic cohort-level variation in the residualized instrument.

Direct cohort-level evidence shows the design has structured identifying variation that the BHJ residualized aggregation cannot detect. Figure 14 displays three panels: pre-residualization cohort means (Panel A) rise monotonically from 0% in 1985 to 93% in 1995, mirroring the LTE rollout; post-residualization \bar{x}_c values (Panel B) lose this monotonic structure because the FE structure has absorbed it; and cohort-by-cohort cross-sectional first-stage slopes (Panel C) become monotonically more negative, from approximately zero in 1985 to a peak magnitude of -6.88 in cohort 1993 (mid-rollout), attenuating to -5.49 in 1995 as saturation sets in.




FIGURE 14. 4G cohort-level identifying variation. Panel A: pre-fixed-effect cohort means of cohort-aggregated LTE coverage (18–25 window mean), 1985–1995 birth cohorts. Panel B: post-fixed-effect residualized \bar{x}_c (the BHJ *ssaggregate* input under county, ACS5-window, and age-bracket fixed effects), shock-level $F = 0.8$. Panel C: cohort-by-cohort cross-sectional first-stage slope $\partial LTE^{18-25} / \partial \log(TRI)$ run separately within each cohort, with 95% CI. *Source:* FCC Form 477 mobile deployment reports 2010–2019; USDA ERS TRI; NHGIS BM6.

Panel C is the textbook signature of a clean rollout-timing instrument: zero before LTE existed, growing in magnitude through the rollout window, peak magnitude at mid-rollout when the largest fraction of counties were transitioning from no-coverage to partial-coverage states, and attenuating toward saturation. The combination of the pre-FE monotonic pattern, the cohort-by-cohort diagnostic, and the conventional county-cluster significance ($t = -5.74$) provides direct evidence of structured identifying variation that the BHJ residualized aggregation cannot detect.

Simulation evidence: BHJ on known-valid designs. To resolve the natural concern that the BHJ failure could reflect weak identification, we generate 500 synthetic cohort-Bartik panels under data-generating processes in which the instrument is mechanically valid by construction, then apply BHJ `ssaggregate` to each. The three DGPs vary only in features that the BHJ failure mode predicts should matter: the *baseline* matches the actual structural features of the 4G sample (3,129 counties, 11 cohorts spanning 1985–1995, sharp post-2010 rollout, ACS5-window \times age fixed-effect structure); the *more-cohorts* variant expands to 31 cohorts (1965–1995, matching the broadband sample range) under the same sharp rollout; and the *smoother-rollout* variant retains 11 cohorts but replaces the sharp rollout with a linear gradient over 2003–2020. In every variant the true 2SLS coefficient is calibrated to the headline estimate in this paper’s primary specification $\beta_{\text{true}} = -0.0034$, errors are calibrated to the residual variance, and identification is mechanically clean by construction.

Table 10 reports the results. The baseline DGP produces median BHJ shock-level $F = 0.88$, with 100 percent of 500 simulation draws falling below the conventional $F = 10$ threshold — essentially identical to the BHJ $F = 0.8$ we obtain in the actual 4G specification. The simulation reproduces our actual BHJ result within rounding on a design that is mechanically valid by construction. The mean 2SLS point estimate recovers β_{true} to four decimal places and median cell-level cluster-robust first-stage F exceeds 300 in every variant, confirming the instrument is valid in all three DGPs. Expanding to 31 cohorts under the same sharp rollout raises median BHJ F only to 5.41, with 100 percent of draws still below threshold: cohort count is not the binding constraint, even at the cohort range of our fixed-broadband specification. Smoothing the rollout to a linear gradient raises median BHJ F to 102.5 with zero percent of draws below threshold. The contrast between the sharp-rollout variants (rows 1–2 of Table 10) and the smooth-rollout variant (row 3) isolates the mechanism: sharp deployment dynamics combined with fixed-effect absorption of time-varying components leave residualized variation on which BHJ aggregation cannot operate. Smoother rollout dynamics survive FE residualization and BHJ recovers.

Three findings emerge from the simulation. First, the Baseline DGP produces BHJ

TABLE 10. BHJ shock-level F in 500 simulations of known-valid cohort-Bartik designs.

DGP variant	Cohorts (N)	Median F_{BHJ}	Share $F < 10$	Mean $\tilde{\beta}$	Median F_{cell}
Baseline (sharp, 11 cohorts)	11	0.88	100%	-0.00339	450
More cohorts (sharp, 31 cohorts)	31	5.41	100%	-0.00340	2,858
Smoother (linear gradient, 11 cohorts)	11	102.5	0%	-0.00340	345
Actual 4G (this paper, Table 8)	11	0.8	—	-0.00342	122

Notes: Each row summarizes 500 simulation draws from a data-generating process in which the instrument is mechanically valid by construction. Each draw generates a synthetic cohort-Bartik panel calibrated to the actual structural features of our 4G specification ($\beta_{\text{true}} = -0.0034$, μ_c, μ_w, μ_a fixed effects matching the actual specification, $\varepsilon_{c,b,a,w} \sim N(0, \sigma^2)$ calibrated to actual residual variance), applies the same 2SLS estimator as the actual specification, and computes BHJ ssaggregate shock-level F . The Baseline DGP matches the structural features of our 1985+ sample. The More-cohorts variant expands the cohort range to 31 (matching the broadband sample) under the same sharp post-2010 rollout dynamics. The Smoother-rollout variant replaces the sharp rollout with a linear gradient over 2003–2020. The final row reports the corresponding values from our actual 4G specification (Table 8) for comparison. Full simulation details, including distribution plots, are in the replication archive.

shock-level $F < 10$ on every one of 500 simulated draws of a known-valid design that matches our structural features, with median $F = 0.88$ matching our actual $F = 0.8$ within rounding. Second, expanding the cohort range to 31 does not rescue the test, even though cell-level F rises proportionally (450 to 2,858): the failure is not driven by small cohort count, which would predict recovery with more cohorts. Third, smoothing the rollout rescues the test entirely while leaving cohort count and identification structure unchanged: the failure is driven by the sharp-rollout \times FE-absorption interaction, exactly the structural feature that distinguishes our specification. We retain the conventional cluster-robust inference reported in Table 8 (cell-level cluster-robust first-stage $F = 122$, well above the conventional threshold), supported by the cohort-by-cohort first-stage diagnostic of Figure 14 and by the simulation evidence in Table 10. We do not interpret this as a general critique of Borusyak, Hull, and Jaravel (2025) ssaggregate, which is the appropriate inference procedure for shift-share designs in many settings; the simulation identifies a specific failure mode — sharp deployment dynamics interacting with fixed-effect absorption of time-varying components — under which the procedure produces sub-threshold F even on designs of known validity. Designs that involve carrier-driven mobile-infrastructure

rollouts or other instruments with similarly sharp temporal dynamics, combined with fixed-effect structures that absorb time-varying components, will face the same issue. We leave the methodological question of how shock-level inference should be conducted in such designs to future work; for the present specification, the cohort-by-cohort first-stage diagnostic and the simulation evidence together support the conventional inference we report.

VII.C Complementary specification: fixed broadband cohort-Bartik

The fixed-broadband cohort-Bartik specification uses the FCC Form 477 residential broadband tier introduced in Section VI, with cohort exposure constructed identically to the LTE specification (mean over the cohort’s 18–25 window; 2008 backfill for pre-2008 years). The cohort treatment intensity variable is the share of the 18–25 window in years 2008+, ranging continuously from $g = 0$ to $g = 1$. The instrument is $z_{c,b}^{\text{BB}} = \log(\text{TRI}_c) \times g_b^{\text{post-2008}}$. In the full 1965–1995 cohort sample, the first stage is strong ($F = 1,018$, first-stage coefficient 0.166, SE 0.005), and the 2SLS estimate is $\beta = -0.0674$ (SE 0.0078, $t = -8.60$). Full broadband adoption (a five-tier change) implies a 0.337 reduction in the ever-married rate, quantitatively similar to the 0.340 implied by the 4G specification at full LTE coverage. The Borusyak, Hull, and Jaravel (2025) shock-level procedure on the 1965–1995 sample yields a shock-level $F = 29$ and a 2SLS coefficient of -0.099 (BHG-robust SE 0.020, $t = -4.97$); the BHJ point estimate is somewhat larger than the cell-level estimate, consistent with observations contributing most to identification also exhibiting larger treatment responses.

Robustness to opioid prescribing exposure. Table 11 reports the broadband 2SLS estimate under a battery of robustness specifications analogous to those in Table 8. Adding cohort-specific opioid prescribing exposure 18–25 as a control (column 2) yields $\beta = -0.054$ (cluster-robust SE 0.011, $t = -5.14$), an attenuation of 19.6 percent from the baseline magnitude. Restricting the sample to non-Appalachian Regional Commission counties (column 3) yields $\beta = -0.064$. Dropping the top quartile of counties by 2008–2014 mean opioid prescribing rate (column 4) yields $\beta = -0.067$, essentially unchanged from the

baseline. Dropping both groups simultaneously (column 5) yields $\beta = -0.064$. The broadband estimate is robust to direct opioid prescribing controls and to geographic exclusions designed to remove the most opioid-affected portions of the sample.

TABLE 11. Fixed broadband cohort-Bartik: robustness to opioid and geographic controls

	(1) Baseline	(2) + opioid 18–25	(3) Drop ARC	(4) Drop top-Q opioid	(5) Drop both
BB ^{18–25}	–0.0674*** (0.0078)	–0.0542*** (0.0105)	–0.0644*** (0.0080)	–0.0673*** (0.0087)	–0.0638*** (0.0085)
First-stage <i>F</i>	1,018	538	906	803	756
Cells	182,972	179,287	156,949	137,010	124,363
Counties	3,129	3,098	2,706	2,373	1,950

Notes: 2SLS estimates of equation (22) with fixed-broadband-tier endogenous variable. Sample: 1965–1995 birth cohorts. Endogenous variable is cohort-aggregated FCC Form 477 broadband tier 18–25 (mean of 8 yearly tiers; backfilled with 2008 value for years pre-2008). Instrument is $\log(\text{TRI}) \times g_b^{\text{post-2008}}$. All specifications include county, ACS5-window, and age-bracket fixed effects. Cells weighted by inverse squared standard error of the cell rate. Standard errors in parentheses are clustered at the county level. Column descriptions parallel Table 8. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The Rotemberg weights diagnostic of Goldsmith-Pinkham, Sorkin, and Swift (2020) shows the broadband identification is well-distributed across the county sample, with the top-five counties holding 1.6 percent of total weight and the top twenty holding 4.9 percent. The Appalachian Regional Commission counties hold 13.3 percent of total weight against a proportional benchmark of 13.6 percent, indicating no concentration on the geography most associated with the opioid epidemic. The top-twenty counties driving identification are predominantly New York, New Jersey, and Oregon metropolitan counties with moderate ruggedness and high broadband adoption, not Appalachian counties with concentrated opioid prescribing.

Cohort-interacted pre-period education. The broadband specification, however, is not robust to the inclusion of cohort-interacted pre-period educational attainment as a control. Adding $\text{bach}_{2000} \times g_b^{\text{post-2008}}$ to the second-stage regression — where bach_{2000} is the Census 2000 bachelor’s-degree-or-higher share at the county level — reduces the first-stage *F* from 1,018 to 1.2 and yields a 2SLS coefficient that flips sign and loses statistical significance.

The first-stage coefficient on the instrument falls 92 percent, from 0.166 to 0.014, with cluster-robust standard error inflating from 0.005 to 0.013. The Frisch-Waugh residualized R^2 of the instrument on the cohort-interacted education control is 0.83: roughly four-fifths of the instrument’s variance is spanned by the cohort-interacted education variable in the full cohort sample.

The mechanism is straightforward. Counties with low Census 2000 bachelor’s-degree share are systematically more rugged than counties with high Census 2000 bachelor’s-degree share, with a cross-sectional correlation between $\log(\text{TRI})$ and bach_{2000} of approximately -0.45 in our sample. The product $\log(\text{TRI}) \times g_b^{\text{post-2008}}$ and the product $\text{bach}_{2000} \times g_b^{\text{post-2008}}$ therefore share most of their identifying variation: both interactions are saying “post-2008 cohorts in rural-rugged-low-education counties experienced different things than post-2008 cohorts in urban-non-rugged-high-education counties.” Adding the second variable as a control absorbs nearly all of the first variable’s identifying signal.

This finding does not arise in the 4G specification of Section VII.B. Adding the analogous $\text{bach}_{2000} \times g_b^{\text{post-2010}}$ control to the 4G IV in the 1985+ cohort sample reduces the first-stage F from 122 to 42 (above the Stock and Yogo (2005) critical value of 16.4 for one endogenous variable at 5 percent bias) but preserves the 2SLS coefficient unchanged at -0.0034 (cluster-robust SE 0.0013, $t = -2.55$). The asymmetry between the broadband and 4G specifications under this control is the primary motivation for treating the 4G specification as the cleaner identification of the digital-infrastructure channel.

The asymmetry reflects a structural difference in deployment economics. Fixed broadband adoption in 2008–2020 was tightly correlated with educational attainment because educated households drove demand for high-speed residential service: where households were willing and able to subscribe, carriers extended service. The cross-sectional pattern of broadband rollout therefore tracks cohort-interacted pre-period education almost mechanically, conditional on terrain. LTE rollout, by contrast, was driven primarily by carrier deployment economics with sharp terrain dependence in cell-tower siting: carriers built towers where the physical infrastructure cost was low and demographic demand was high enough to justify the deployment, but tower siting depends more on terrain feasibility (line

of sight, slope grade, right-of-way acquisition) than on local educational composition. The result is that 4G rollout has identifying variation in the terrain instrument that is structurally distinct from the variation captured by cohort-interacted pre-period education.

Bias from correlated rural-mortality channels and joint-bundle interpretation. The broadband estimate is also exposed to deaths-of-despair-correlated rural-mortality bias. Table 9 reports implied bias of 71–120 percent of headline magnitude under NCHS broad-drug mortality across the FE-only and progressively-controlled specifications. The bias is specific to the broader NCHS measure: under WONDER opioid-specific mortality, broadband implied bias is essentially zero, so the bias arises through mortality channels broader than opioids — cocaine, methamphetamine, sedative, and unspecified drug poisonings, which together capture the deaths-of-despair pattern documented by Case and Deaton (2017, 2020) but exceed the prescription-opioid epidemic in scope.

Taken together, the cohort-interacted education absorption and the NCHS broad-drug bias arithmetic indicate that the broadband cohort-Bartik specification does not separately identify the digital-infrastructure channel from co-located rural-trajectory factors. The defensible interpretation of the broadband 2SLS estimate is a joint effect of post-2008 differential trajectories in rugged-rural counties — a bundle that includes digital-infrastructure rollout, deaths-of-despair-correlated rural mortality, and persistent rural-education-decline factors. We retain the broadband specification for two reasons: the magnitude convergence with the 4G specification (-0.337 versus -0.340 at full treatment) provides cross-specification corroboration of effect size, and the broadband sample includes the 1965–1985 cohort range that the smartphone-era restriction excludes, corroborating that the effect is not specific to smartphone-era cohorts but operates across the longer arc of digital adoption since the late 1990s.

VII.D Bounded inference and implications for the structural model

Combining the two specifications, we report the digital-infrastructure effect on the cohort ever-married rate as a bounded estimate. The 4G specification on the 1985–1995 smartphone-era cohorts gives a point estimate of -0.340 at full LTE coverage with bounded

bias between 0 and 8 percent of headline. The fixed-broadband specification on the 1965–1995 cohort range gives a joint bundled point estimate of -0.337 at full broadband adoption; the broadband-only contribution is bounded by the cohort-interacted education absorption and the rural-mortality bias arithmetic at 30–80 percent of the joint estimate. The two specifications support an effect on cohort partnership formation in the range -0.30 to -0.34 at full digital saturation, with magnitude convergence across structurally distinct identifying variation providing the strongest joint claim. The within-country identification of Sections VI and VII is one of four evidence sources for the digital channel; the others are the cross-country falsification of Section V.G, the causal teen-fertility identification of Section VI.A, and the structural model fits of Section V.B.

The 4G specification’s cohort range (1985–1995) corresponds precisely to the theoretically-relevant population that Section III identifies: cohorts whose 18–25 windows overlap the post-2007 cell-phone-network tipping point and the post-2010 LTE rollout. The fixed-broadband specification’s wider cohort coverage corroborates effect magnitude but identifies a joint bundle that includes pre-smartphone-era digital-infrastructure rollout, consistent with the structural model’s treatment of the digital shock as a continuous-decline price trajectory beginning before 2007 rather than a discrete 2007 shock.

At the 4G specification’s identified -0.34 , the digital-infrastructure channel accounts for a substantial fraction of the cohort-level marriage decline among the 1985+ smartphone-era cohorts. The 1985 cohort completed age 25 with an ever-married rate of approximately 0.56 in the ACS cells; the 1995 cohort approached comparable age points at approximately 0.35, a decline of ~ 0.21 . The 4G effect at the cohort-mean LTE coverage differential (~ 50 percentage points between the 1985 and 1995 cohorts’ 18–25 windows, reflecting that the 1995 cohort’s window captured only the back half of the post-2010 rollout) accounts for ~ 0.17 of the cohort-mean decline, or roughly 81 percent of the observed extensive-margin partnership-formation collapse on this cohort range. The remaining 19 percent corresponds to factors operating outside the digital-infrastructure channel — the secular decline in marriage rates documented since the 1970s, post-2008 economic-trajectory factors, and idiosyncratic cohort variation in marriage timing — consistent with the paper’s central

framing that smartphones accelerated a partnership-formation decline that was already in progress for reasons documented in prior literature, rather than initiating it ab initio.

VIII. CONCLUSION

The post-2007 collapse of fertility in high-income countries reflects a Becker-style quantity-quality reallocation on *connections* rather than children, induced by digital technology that lowered the relative price of broad-shallow ties while leaving deep-relationship production unchanged. Time reallocated from depth-producing activities to digital leisure and shallow-connection maintenance; the depth stock that supports partnership formation eroded; partnership formation collapsed on the extensive margin and conditional fertility fell on the intensive margin. Calibrated to US time-use and fertility data over 2007–2024, the model accounts for essentially the full US fertility decline (101 percent), with a counterfactual that holds digital quality at its 2007 level explaining 43 percent of the observed decline. Applied with country-specific wage paths to the full panel of 38 OECD member countries, the median observed and median model-predicted TFR declines align to within one percentage point on the 2007–2023 window, with 84 percent (32 of 38) of countries showing the model-predicted direction of decline. The cumulative OECD-aggregate fit through 2023 lines up to one percentage point (model –18.5 percent, data –18.4 percent). Countries where the model gets the sign wrong are concentrated among countries with partial post-2007 recoveries from very low pre-2007 fertility (Czech Republic, Hungary, Slovakia, Slovenia, and Portugal) and Germany, where family-policy recovery moved fertility in the opposite direction from the model’s monotone shock; we report all 38 cases transparently in the figure.

The mechanism is causally identified at the partnership-formation margin: at the youngest end through the IV of Section VI.A, and across the smartphone-era cohorts through the cohort-Bartik design of Section VII. At the youngest end, Section VI.A estimates that 10 additional percentage points of long-run county-level broadband caused an 18.9 percentage point larger decline in teen birth rates over 2003–2018 ($F = 73$). At

the partnership margin itself, Section VII uses a cohort-Bartik instrument on the full 3,129-county US sample of female ever-married rates: the primary 4G LTE specification on the 1985–1995 smartphone-era cohorts gives a -0.34 effect at full LTE saturation ($F = 122$, robust to opioid, geographic, and cohort-interacted pre-period-education controls; bias arithmetic at 0–8 percent of headline). A complementary fixed-broadband specification on the wider 1965–1995 cohort range corroborates the magnitude through structurally distinct identifying variation; together the two specifications place the effect between -0.30 and -0.34 , accounting for roughly 80 percent of the cohort-mean ever-married decline among smartphone-era cohorts. Section VI extends the cohort-exposure logic to age-specific fertility, recovering the predicted age gradient with a placebo-passing oldest cohort.

The model is best read as a structural extension of unified growth theory rather than a departure from it. The first demographic transition was a quantity-quality reallocation on children, induced by industrialization that raised the return to child quality; the second is a quantity-quality reallocation on connections, induced by digitalization that lowered the relative cost of breadth. The same logic operates in both, on a different factor in household production.

REFERENCES

- Aguiar, Mark, Mark Bilal, Kerwin Charles, and Erik Hurst. “Leisure Luxuries and the Labor Supply of Young Men.” *Journal of Political Economy* 129, no. 2 (2021): 337–382.
- Aizcorbe, Ana M., 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.” *NBER Working Paper 25645 / FEDS Working Paper 2019-012*, 2019.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. “The Skill Complementarity of Broadband Internet.” *Quarterly Journal of Economics* 130, no. 4 (2015): 1781–1824.
- Allred, Colette. “Refined Divorce Rate in the U.S.: Geographic Variation, 2024.” NCFMR Family Profile FP-25-31. Bowling Green, OH: National Center for Family & Marriage Research, 2025.
- Allred, Colette. “Age Variation in the Refined Divorce Rate, 1990 & 2023.” NCFMR Family Profile FP-25-24. Bowling Green, OH: National Center for Family & Marriage Research, 2025.
- Büyükeren, Berkeren, Alexey Makarin, and Heyu Xiong. “The Impact of Dating Apps on Young Adults: Evidence from Tinder.” *American Economic Journal: Applied Economics* 18, no. 2 (2026): 61–106.
- Becker, Gary S. “An Economic Analysis of Fertility.” In *Demographic and Economic Change in Developed Countries*. Princeton: Princeton University Press, 1960.
- Becker, Gary S., and H. Gregg Lewis. “On the Interaction Between the Quantity and Quality of Children.” *Journal of Political Economy* 81, no. 2 (1973): S279–S288.
- Becker, Gary S. *A Treatise on the Family*. Cambridge, MA: Harvard University Press, 1981.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel. “A Practical Guide to Shift-Share Instruments.” *Journal of Economic Perspectives* 39, no. 1 (2025): 181–204.

- Bradbury, Thomas N., and Benjamin R. Karney. "Understanding and Altering the Longitudinal Course of Marriage." *Journal of Marriage and Family* 66, no. 4 (2004): 862–879.
- Brown, Susan L., and I-Fen Lin. "The Graying of Divorce: A Half Century of Change." *Journals of Gerontology Series B* 77, no. 9 (2022): 1710–1720.
- Brown, Susan L., I-Fen Lin, and Kagan A. Mellencamp. "The Rising Midlife First Marriage Rate in the U.S." *Journals of Gerontology Series B* 78, no. 7 (2023): 1141–1150.
- Byrne, David M. "The Mysterious Cross-Country Dispersion in Mobile Phone Price Trends." *FEDS Notes*, August 5, 2019. Washington: Board of Governors of the Federal Reserve System.
- Cacioppo, John T., and William Patrick. *Loneliness: Human Nature and the Need for Social Connection*. New York: W. W. Norton, 2008.
- Carroll, Christopher D. "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems." *Economics Letters* 91, no. 3 (2006): 312–320.
- Case, Anne, and Angus Deaton. "Mortality and Morbidity in the 21st Century." *Brookings Papers on Economic Activity* (Spring 2017): 397–476.
- Case, Anne, and Angus Deaton. *Deaths of Despair and the Future of Capitalism*. Princeton: Princeton University Press, 2020.
- Caudillo, Mónica L., Andrés Villarreal, and Philip N. Cohen. "The Opioid Epidemic and Family Formation in the United States." Working paper, University of Maryland, 2022.
- Centers for Disease Control and Prevention, National Center for Health Statistics. *Nativity, 2007–2024 Expanded*, CDC WONDER online database. Available at <https://wonder.cdc.gov/nativity-current.html>, accessed 2026.
- Cigna. *Loneliness in America: 2020 U.S. Loneliness Index*. Bloomfield, CT: Cigna Corporation, 2020.
- Cohen, Sheldon, and Thomas A. Wills. "Stress, Social Support, and the Buffering Hypothesis." *Psychological Bulletin* 98, no. 2 (1985): 310–357.
- 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.
- Dunbar, R. I. M. "The Anatomy of Friendship." *Trends in Cognitive Sciences* 22, no. 1 (2018): 32–51.
- Eijkemans, Marinus J. C., Frans van Poppel, Dik F. Habbema, Ken R. Smith, Henri Leridon, and Egbert R. te Velde. "Too Old to Have Children? Lessons from Natural Fertility Populations." *Human Reproduction* 29, no. 6 (2014): 1304–1312.
- Eurostat. Harmonised European Time Use Survey, various waves.
- Facebook (Meta Platforms, Inc.). "One Billion People on Facebook." Press release, October 4, 2012, <https://about.fb.com/news/2012/10/one-billion-people-on-facebook/>.
- Galor, Oded, and David N. Weil. "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond." *American Economic Review* 90, no. 4 (2000): 806–828.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110, no. 8 (2020): 2586–2624.
- Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. "Engines of Liberation." *Review of Economic Studies* 72, no. 1 (2005): 109–133.
- Haidt, Jonathan. *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. New York: Penguin Press, 2024.
- Hellstrand, Julia, Jessica Nisén, and Mikko Myrskylä. "All-Time Low Period Fertility in Finland: Demographic Drivers, Tempo Effects, and Cohort Implications." *European Journal of Population* 38, no. 2 (2022): 191–221.
- Holt-Lunstad, Julianne, Timothy B. Smith, and J. Bradley Layton. "Social Relationships and Mortality Risk: A Meta-analytic Review." *PLOS Medicine* 7, no. 7 (2010): e1000316.
- House, James S. *Work Stress and Social Support*. Reading, MA: Addison-Wesley, 1981.
- Hudson, Nathan, and Hernan J. Moscoso Boedo. "The Collapse of Teen Fertility in the Digital Era." April 29, 2026. Available at SSRN: <https://ssrn.com/abstract=6676839>.
- Kearney, Melissa S., and Phillip B. Levine. "Why Is Fertility So Low in High Income Countries?" NBER Working Paper 33989, revised January 2026.

- 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.
- Khan, Diba, Brady 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. DOI: <https://dx.doi.org/10.15620/cdc/20250305002>.
- Krämer, Nicole C., Vera Sauer, and Nicole Ellison. "The Strength of Weak Ties Revisited: Further Evidence of the Role of Strong Ties in the Provision of Online Social Support." *Social Media + Society* 7, no. 2 (2021): 1–12.
- Loo, Jaden. "Divorce: More than a Century of Change, 1900–2022." NCFMR Family Profile FP-24-11. Bowling Green, OH: National Center for Family & Marriage Research, 2024.
- Mahler, Lukas, Michèle Tertilt, and Minchul Yum. "Policy Concerns in an Era of Low Fertility: The Role of Social Comparisons and Intensive Parenting." *Brookings Papers on Economic Activity*, Fall 2025; CEPR Discussion Paper 20767, December 2025.
- McPherson, Miller, Lynn Smith-Lovin, and Matthew E. Brashears. "Social Isolation in America: Changes in Core Discussion Networks Over Two Decades." *American Sociological Review* 71, no. 3 (2006): 353–375.
- Murthy, Vivek H. *Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community*. Washington, DC: U.S. Department of Health and Human Services, 2023.
- Newzoo. *Global Mobile Market Report 2021*. Amsterdam: Newzoo, 2021.
- Nunn, Nathan, and Diego Puga. "Ruggedness: The Blessing of Bad Geography in Africa." *Review of Economics and Statistics* 94, no. 1 (2012): 20–36.
- Organisation for Economic Co-operation and Development. *OECD Family Database, SF2.1 Fertility Rates*. Paris: OECD, 2024. Available at <https://www.oecd.org/els/family/database.htm>.
- Orben, Amy, and Andrew K. Przybylski. "The Association Between Adolescent Well-being and Digital Technology Use." *Nature Human Behaviour* 3, no. 2 (2019): 173–182.
- Pew Research Center. "A Record-High Share of 40-Year-Olds in the U.S. Have Never Been Married." Pew Research Center Social and Demographic Trends, June 2023.
- Pew Research Center. "Share of U.S. Adults Living Without a Romantic Partner Has Ticked Down in Recent Years." Pew Research Center Short Reads, January 2025.
- Primack, Brian A., Ariel Shensa, Jaime E. Sidani, Erin O. Whaite, Liu yi Lin, Daniel Rosen, Jason B. Colditz, Ana Radovic, and Elizabeth Miller. "Social Media Use and Perceived Social Isolation Among Young Adults in the U.S." *American Journal of Preventive Medicine* 53, no. 1 (2017): 1–8.
- Putnam, Robert D. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster, 2000.
- Riley, Shawn J., Stephen D. DeGloria, and Robert Elliot. "A Terrain Ruggedness Index that Quantifies Topographic Heterogeneity." *Intermountain Journal of Sciences* 5, no. 1–4 (1999): 23–27.
- Roberts, Sam G. B., and Robin I. M. Dunbar. "The Costs of Family and Friends: An 18-Month Longitudinal Study of Relationship Maintenance and Decay." *Evolution and Human Behavior* 32, no. 3 (2011): 186–197.
- Rosenfeld, Michael J., Reuben J. Thomas, and Sonia Hausen. "Disintermediating Your Friends: How Online Dating in the United States Displaces Other Ways of Meeting." *Proceedings of the National Academy of Sciences* 116, no. 36 (2019): 17753–17758.
- Rosenfeld, Michael J., and Sonia Hausen. *How Couples Meet and Stay Together (HCMST) 2017, 2020, 2022, United States*. Inter-university Consortium for Political and Social Research, ICPSR 38873, 2024.
- Statista. *Global Smartphone Penetration Rate by Region and Country, 2016–2023*. Hamburg: Statista, 2024.
- Steiner, Anne Z., and Anne Marie Z. Jukic. "Impact of Female Age and Nulligravidity on Fecundity in an Older Reproductive Age Cohort." *Fertility and Sterility* 105, no. 6 (2016): 1584–1588.
- Stock, James H., and Motohiro Yogo. "Testing for Weak Instruments in Linear IV Regression." In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, edited by Donald W. K. Andrews and James H. Stock, 80–108. Cambridge: Cambridge University Press, 2005.
- The Knot. *The Knot 2024 Real Weddings Study*. The Knot Worldwide, 2024.

- Turkle, Sherry. *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books, 2011.
- Turkle, Sherry. *Reclaiming Conversation: The Power of Talk in a Digital Age*. New York: Penguin Press, 2015.
- 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. Census Bureau. *Current Population Survey, Annual Social and Economic Supplement, Marital Status Tables*. Washington, DC: U.S. Census Bureau, 2023.
- Westrick-Payne, Karen Kasey. "Refined Marriage Rate in the U.S.: Geographic Variation, 2024." NCFMR Family Profile FP-25-30. Bowling Green, OH: National Center for Family & Marriage Research, 2025.
- World Bank. *World Development Indicators: Fertility Rate, Total (Births per Woman), Series SP.DYN.TFRT.IN*. Washington, DC: World Bank, 2024. Available at <https://data.worldbank.org/indicator/SP.DYN.TFRT.IN>.
- Yanagimoto, Kazuharu. "A Quantitative Model of Non-Marriage and Fertility: Bargaining over Leisure." Working paper, Kobe University and CEMFI, August 2024 (revised March 2026).
- Yoo, Sam Hyun. "Later, Fewer, None? Recent Trends in Cohort Fertility in South Korea." *Demography* 60, no. 2 (2023): 563–582.