DISENGAGED: DIGITAL LIFE, MARRIAGE, AND THE YOUNG MALE LABOR-FORCE DECLINE

NATHAN HUDSON AND HERNÁN MOSCOSO BOEDO[†]

July 21, 2026

ABSTRACT

Labor-force participation among young non-college educated men fell 3.2 points after the mid-2000s and, alone among major groups, never recovered. Part of the decline ran through marriage rather than work. Single men's participation rate never broke; what shifted was composition. Because married men participate eight to eleven points above single men, the roughly ten-point fall in the married share explains about a third of the decline. Run identically for women, the same shift reverses sign and explains most of a symmetric rise, redistributing labor supply across the genders. A within-zone cohort comparison ties the decline to the 4G rollout.

Keywords: labor force participation, mobile internet, smartphones, time use, marriage, young men.

JEL Codes: J12, J21, J22, O33, L96.

I. INTRODUCTION

Labor-force participation among young non-college educated men in the United States declined steadily over the two decades before the COVID-19 pandemic and, alone among major groups, failed to recover after the Great Recession. The fact has anchored a large literature on the changing position of less-educated men (Eberstadt, 2016; Krueger, 2017; Binder and Bound, 2019; Abraham and Kearney, 2020), and the candidate explanations (collapsing manufacturing demand, opioids and disability, the rising quality

[†]Department of Economics, Carl H. Lindner College of Business, University of Cincinnati. Corresponding author: Hernán Moscoso Boedo, hernan.moscoso@uc.edu. No external funding supported this research; the authors report no conflicts of interest. All data are public (IPUMS ACS and ATUS, FCC mobile-deployment filings); replication code is available from the authors.

of recreational computing) are numerous and not mutually exclusive. This paper does not displace them. It isolates a part of the decline they have not jointly confronted and shows that its first-order channel was not work but marriage.

The argument rests on a decomposition. Aggregate participation among young non-college educated men can be written as a population-weighted average of participation within partnership cells, $L_t = \sum_g s_{g,t}\,\ell_{g,t}$, and its change over time splits into a within component (cell participation rates moving) and a between component (the population reallocating across cells). Throughout, we call the first object the *within-cell rate*, the participation rate of a group such as single non-college educated men taken on its own, and the second *composition*, the shares of the population sitting in the married and single cells. The paper's claim, in these terms, is that the aggregate moved mainly because the shares moved, not because either group's own rate bent at the smartphone's arrival. When we carry out this decomposition on the American Community Survey, two facts emerge. First, the participation rate of single non-college educated men shows no break at the smartphone era: it was already falling before 2007 and continued at essentially the same pace afterward, so the within component is the continuation of a pre-existing trend. Second, what changed was composition. The share of young non-college educated men who are married fell about ten percentage points between the 2005–07 and 2015–19 windows, continuing a marriage decline that long predates the smartphone and proceeds smoothly in calendar time. Because married men participate eight to eleven points above single men, this reallocation accounts for roughly a third of the aggregate decline. We claim no break in the calendar path of either series; the digital content of the argument rests on the cohort comparison of Section VI, where marriage falls for precisely the cohorts whose formative years overlapped the rollout, within place. The two facts do not compete: the digital era's measurable imprint on these men's aggregate labor supply operates through who they are, not through any visible bending of single men's own participation rate.

The same marriage-market event yields what we take to be the descriptive record's most striking fact. Carrying out the identical decomposition for young non-college educated women, the married-share decline enters with the opposite sign, because single women participate above their married counterparts rather than below. Aggregate participation among young non-college educated women rose 3.2 points over the same window in which the men's fell 3.2, and the larger part of that rise is the same marriage-market shift running in reverse. The single shares of young non-college educated men and women move together to within a point, as the two sides of one market must. The digital-era marriage decline did not destroy young non-college labor supply so much as redistribute it across the genders. Time-use data sharpen the reversal: among young non-college educated women, market work is higher for the single, but total work is not: once home production is accounted for, the married supply more in total, as married men do (Section III). This symmetry also fixes the paper's scope: the object that requires a behavioral explanation is male, because both of the women's cells rose within-cell while among men only the single cell fell appreciably; the married cell was essentially flat.

If marriage is the channel, the central question is what moved marriage. The behavioral record points to the reorganization of social life. Using ATUS time diaries with companionship codes, we show that young non-college educated single men reallocated thirty-four minutes per day into solitary screen leisure (games, computer use for leisure, and digitally mediated communication recorded with no other person present) between the pre- and post-smartphone windows, against a change of five minutes for college-educated married men, the anchored comparison group. Over the same windows these men withdrew twenty minutes per day from time spent with friends. A full accounting of the day locates where the screen time came from: it came from other leisure and from time with others, while market work held flat, so the rising value of solitary leisure that the leisure-luxury account invokes did not, in this cell, draw time from work. The reallocation of social time is an order of magnitude larger than the reallocation of labor supply, and that asymmetry disciplines what follows: whatever mobile internet did to these men, it did far more to their social lives than to their hours of work, which is exactly the pattern a marriage channel predicts and a direct labor-supply channel does not.

We organize these facts with a household-production model and then ask what drove the marriage decline at their center. The model, based on Hudson and Moscoso Boedo

(2026b), with an endogenous labor margin, delivers the married–single participation gap from a single primitive: because a deeper relationship supports a larger household and the digital bundle and time are gross complements, providing for that household raises the marginal value of market work, so relational capital and labor supply move together. Calibrated to the falling smartphone hedonic price, with nothing fit to labor, the model reproduces the fertility decline as an untargeted prediction and, through composition, traces a participation decline of the size the decomposition measures. We then test the digital origin directly. The habits and relationships that lead to marriage take shape in adolescence and early adulthood, and mobile-broadband saturation arrived in different places in different years, so birth cohorts within the same commuting zone differ in how much of their formative years fell under local coverage. Comparing cohorts within zone, which differences out the urban–rural gradient that confounds cross-place comparisons, the cohorts who came of age on 4G married less, in a pattern that is flat across the pre-rollout cohorts and breaks at exactly the exposed ones, tracks the timing of the rollout rather than of the Great Recession, and appears in parallel for women. We are deliberate about what this shows: the break is concentrated among non-college educated men, the group whose participation fell, and absent for college men, while it appears for women of both schooling levels. We read the cohort evidence as corroboration of a digital-timed marriage response and rest the paper's weight on the composition accounting and the cross-gender redistribution, which require no instrument.

The paper therefore makes two claims held to two different standards of evidence, and it may help to keep the hierarchy explicit. The first is an accounting claim, and it is exact: given the observed fall in the married share, the married–single participation gap converts that fall into roughly a third of the men's aggregate decline, and the same arithmetic run for women delivers most of their rise. The second is a causal claim about why marriage fell, and it is held to a corroboration standard: the cohort evidence shows the marriage decline arriving on the timing, in the places, and in the populations a digital origin would produce, and we present it as consistent with that origin rather than as a stand-alone causal magnitude. Nothing in the first claim depends on the second, and readers whose primary interest is the causal question may proceed directly to Section VI.

The paper contributes to three literatures. To the literature on the labor supply of young men (Aguiar et al., 2021; Krueger, 2017; Binder and Bound, 2019; Abraham and Kearney, 2020), it contributes the composition result, that roughly a third of the aggregate decline among young non-college educated men runs through the partnership margin rather than within-group behavior, together with time-diary evidence that the screen leisure these men gained came at the expense of other leisure and of time with friends rather than of work. To the literature on internet access and labor markets, which finds zero or positive employment effects of connectivity in developed economies (Dettling, 2017; Bhuller et al., 2020; Atasoy, 2013; Zuo, 2021) and positive effects in developing ones (Hjort and Poulsen, 2019), it contributes a mechanism that work focused on search, matching, and skill complementarity has not isolated: a decline in partnership formation that reaches aggregate labor supply through demographic composition, with opposite-signed consequences for men and women. The contribution closest to ours, Myers and Hooper (2026), uses the iPhone–AT&T exclusivity window to estimate effects of smartphone access on fertility; we study a distinct margin, the formation of marriages and its transmission to labor supply, building on the partnership and fertility evidence of Hudson and Moscoso Boedo (2026a, 2026b).

The remainder of the paper proceeds as follows. Section II places the argument in the related literature. Section III documents the decline, its composition, and the cross-gender redistribution; Section IV traces the reallocation of time in the ATUS. Section V develops the model of relational capital and labor supply and calibrates it to the smartphone price path. Section VI presents the cohort-exposure evidence on the marriage decline's digital timing, and Section VII concludes.

II. RELATED LITERATURE

The decline in participation and hours among young non-college educated men is among the most studied facts in modern U.S. labor economics. Aguiar et al. (2021) attribute a portion of the decline to improvements in the quality of recreational computing, which raise the value of leisure relative to work. Krueger (2017) emphasizes pain, disability, and opioid use; Binder and Bound (2019) emphasize the joint deterioration of these men's position in the marriage market and the labor market; Eberstadt (2016) documents the scale of male labor-force exit; Abraham and Kearney (2020) review the full menu of demand- and supply-side explanations and conclude that no single factor dominates; Autor et al. (2013) trace much of the regional decline to collapsing manufacturing demand. These explanations divide on what our evidence can say to them. The demand-side and health channels operate on the single cell's own attachment to work, and our decomposition leaves them their room: the within-cell participation rate is real and roughly two-thirds of the aggregate decline, and we do not displace it. The leisure-luxury channel makes a sharper and more testable claim, that better screen leisure substitutes for work within the individual, and it is the one our time-use evidence speaks to directly. In the cell where screen time rose most, we find the market-work margin unmoved and the screen increase drawn from other leisure, not from work (Section IV); the aggregate hours decline that channel documents is, on this evidence, substantially the same compositional shift that moves participation, married men working more than single men as the married share falls. Our contribution is to isolate that compositional third and to point to a cause for it.

That married men work and participate more than single men is one of the most durable regularities in labor economics, and the eight-to-eleven-point gap we measure within young non-college educated men is consistent with a literature going back decades (Korenman and Neumark, 1991; Ahituv and Lerman, 2007). The interpretation of the gap is contested: it reflects some mixture of selection, in that more employable men marry, and a causal effect of marriage on labor supply through specialization and the demands of providing, and we are deliberately agnostic about the split. Our chain does not require marriage to cause participation; it requires only that the two are associated at the rate we measure, so that a change in the share married mechanically shifts aggregate participation through composition. What the cohort comparison supplies that this literature has generally lacked is plausibly exogenous variation in marriage itself: most estimates of the marriage–labor relationship struggle to separate selection from effect precisely because marriage is chosen, and we sidestep the problem by studying what moved whether these cohorts married and letting the established gap do the accounting.

The empirical literature on internet access and employment in developed economies finds, with striking consistency, effects that are zero or positive. Broadband raised married women's participation (Dettling, 2017), improved labor-market matching (Bhuller et al., 2020), and raised employment where it complemented skill (Akerman et al., 2015; Atasoy, 2013); subsidized residential broadband raised low-income employment (Zuo, 2021); in developing economies fast internet raised employment substantially (Hjort and Poulsen, 2019). Against this backdrop a negative effect of mobile internet on young men's attachment to work would be anomalous if it operated through the channels this literature studies. Our results suggest the channel is different in kind: not a labor-market channel at all but a family-formation channel, a decline in the rate at which young men formed the partnerships that, through composition, sustain attachment to work. This connects the labor literature to a parallel body of work on digital technology and social life: the displacement of in-person socializing by screen time (Aguiar et al., 2021; Kimbrough, 2020), the substitution of digital for face-to-face couple formation (Rosenfeld et al., 2019), and the partnership-formation and fertility consequences developed in our prior work (Hudson and Moscoso Boedo, 2026a,b). The channel also connects to the causal literature on social media and wellbeing, in which the staggered rollout of Facebook (Braghieri et al., 2022) and the experimental deactivation of Allcott et al. (2020) identify effects of digital connection operating through displaced in-person contact. Our contribution is to trace an analogous mechanism to a consequential new margin, the formation of marriages, and through it to labor supply.

III. FACTS: THE DECLINE AND ITS COMPOSITION

This section sets out the empirical facts the rest of the paper explains. We first describe the data and the cells; we then document the post-2007 decline in young non-college educated men's labor-force participation, show that it is compositional rather than a within-cell change, and trace its mirror image in the labor supply of young women.

III.A Data and cells

The labor-force analysis draws on the American Community Survey (2005–2019) via IPUMS (Ruggles et al., 2024). The population of interest is men and women aged 21–30 who are not enrolled in school, partitioned by education (no four-year degree versus four-year degree or more) and partnership status. We refer to the four cells as NCS (non-college single), NCM (non-college married), CS (college single), and CM (college married), with NCS the population of interest and CM the anchored comparison group. The ACS identifies legal marital status but not cohabitation; throughout, "married" denotes a current legal marriage (spouse present or absent) and "single" its complement. We verified that the composition and decomposition results below are insensitive to this distinction, reproducing them to the decimal under the broader married-or-cohabiting definition where it can be constructed. The time-use analysis of Section IV uses the American Time Use Survey, 2003–2023, via IPUMS (Flood, Sayer, and Backman, 2023), with the activity and companionship ("who") files; the cohort comparison of Section VI adds commuting-zone panels of 4G coverage built from FCC mobile-deployment filings (Federal Communications Commission, 2020). Variable definitions specific to each exercise are given where it is introduced. Table 1 reports summary statistics for the three analysis samples.

III.B The within-cell rate did not break; composition did

Table 2 reports the partnership composition and cell participation of young men and women by education; we read the non-college educated men here and the women's panel below. Between 2005–07 and 2015–19 their married share fell 10.6 percentage points within the non-college group, the single cell absorbing the difference one-for-one. The within-cell participation rate of single non-college educated men, by contrast, shows no smartphone-era break: it stood at 84.9 percent in 2005–07 and 82.0 percent in 2015–19, a continuation of a slide already underway before 2007, while the married cell's rate was

TABLE 1. Summary statistics, analysis samples

Panel A. ACS labor sample: non-college, aged 21–30, not enrolled

	Men		Women	
	2005–07	2015–19	2005–07	2015–19
Observations	305,500	555,000	260,700	434,600
Married (%)	31.7	21.1	41.4	30.0
In labor force (%)	87.52	84.35	71.55	74.76
Mean age	25.6	25.7	25.7	25.8

Panel B. Cohort event-study sample: ever-married, aged 25–34, all education, cohorts 1983–96

	Men	Women
Commuting-zone × cohort cells	6,436	6,360
Commuting zones	493	492
Birth cohorts	14	14
Respondents (persons)	2,097,010	2,046,918
Median cell size	123	116
Ever married (%)	38.7	49.4
CZ 4G coverage, 2013 (cell mean; s.d.)	0.29 (0.32)	0.29 (0.32)

Panel C. ATUS time-use sample: non-college educated single men, aged 21–30

	2003–07	2015–19
Respondent-days	1,446	959

Notes. The samples differ by design: the labor and time-use analyses (Panels A and C) study non-college educated men and women aged 21–30, the population whose participation fell, while the cohort marriage design (Panel B) uses ages 25–34, over which marriage is observable, and pools all education for power, with the education breakdown in Section VI.F. Panel A: ACS non-college educated men and women aged 21–30, not enrolled, by period; every statistic is computed on the non-college cell; observation counts are unweighted respondent counts rounded to the nearest hundred, and the married share, labor-force rate, and mean age are person-weighted. Panel B: the estimating sample of the cohort event study (Section VI): commuting-zone × birth-cohort cells with at least 25 sampled respondents, together with the 2013 commuting-zone 4G coverage that supplies the treatment (its mean and s.d. taken across cells). Panel B cohort counts and the median cell size are per gender (men 493 commuting zones, women 492). Panel C: the single non-college educated male cell of the daily time budget (Table 6). Sources: ACS and ATUS via IPUMS; FCC mobile-deployment filings.

essentially flat at roughly 93 percent. The compositional margin is where the aggregate change concentrates; its own calendar path is steady rather than sharp (the married share was falling before 2007 and continued falling after), and the digital timing the paper attributes to it is established by the cohort comparison of Section VI, not by the calendar record. The within-cell rate is the province of the longer-run explanations we do not displace.

An annual reading of both margins confirms this division of labor. Using the ACS year by year (2005–2023) with a 2000 Census anchor on the married share (the cell-rate fits use 2005–2023 alone, Census-era labor-force coding being incomparable across the 2000–2005 gap), the cell rates carry no smartphone-era break: a piecewise-linear fit to the single non-college cell's participation rate rejects neither a slope change at 2007 ($t = -1.30$) nor a level break ($t = -1.21$), and the married cell holds near 93 percent throughout. The married share itself declines across the whole sample, and declined *faster* before 2007 (about -1.7 points per year) than after (about -0.6), flattening after 2015; in proportional terms the deceleration is milder (roughly -4.0 versus -3.0 log points per year), and part of the level movement around 2008 reflects recession-era school enrollment changing selection into the not-enrolled cell rather than marriage itself. The calendar record therefore contains no 2007 discontinuity in either object, and we claim none: the ten-point window decline that drives the accounting is an endpoint fact, true whatever the path's shape, and the digital timing of the marriage decline is identified in Section VI by cohort formative exposure within place, which is where, under the mechanism, it should live. Figure 1 plots both series.

III.C A within/between decomposition

To quantify how much of the aggregate decline operates through composition, write aggregate participation among young non-college educated men as $L_t = \sum_g s_{g,t}\,\ell_{g,t}$, with $g \in \{\text{single}, \text{married}\}$. The change between periods decomposes as

$$\Delta L \;=\; \underbrace{\sum_g s_{g,0}\,\Delta\ell_g}_{\text{within}} + \underbrace{\sum_g \ell_{g,1}\,\Delta s_g}_{\text{between}} \;=\; \underbrace{\sum_g s_{g,1}\,\Delta\ell_g}_{\text{within}} + \underbrace{\sum_g \ell_{g,0}\,\Delta s_g}_{\text{between}}, \tag{1}$$

two algebraically exact decompositions: the first weights the within term by base-period shares $s_{g,0}$ and the between term by end-period rates $\ell_{g,1}$ (a Laspeyres reading), the second reverses both, weighting the within term by end-period shares and the between term by base-period rates (a Paasche reading). Each is exact; the share–rate interaction is absorbed

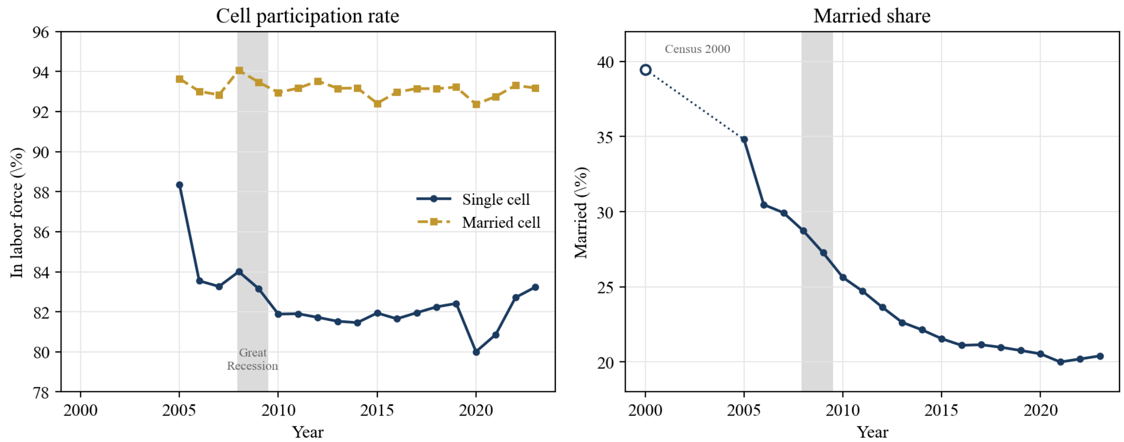

FIGURE 1. Annual paths of the cell participation rates and the married share, non-college educated men aged 21–30, ACS 2005–2023 with a 2000 Census anchor. Panel A: the single and married cells' own participation rates; neither breaks at the smartphone's arrival (single cell: slope-change $t = -1.30$, level-break $t = -1.21$), the single cell continuing a pre-existing slide and the married cell holding near 93 percent. Panel B: the married share declines throughout, faster before 2007 than after, and flattens after 2015. The shaded band in each panel marks the Great Recession. The decomposition of Section III uses only the 2005–07 and 2015–19 endpoint values of these series, a difference that is invariant to the path between them.

differently in the two, so they bracket the within and between contributions, and we report the range across them. Aggregate participation among young non-college educated men fell 3.17 points, from 87.52 to 84.35 percent (Table 3). The within component contributes -2.01 to -2.29 points (64 to 72 percent of the total); the between component, the reallocation from the high-participation married cell to the lower-participation single cell, contributes -0.87 to -1.15 points (28 to 36 percent). Composition is not the majority of the decline, but it is the only component for which the paper offers a candidate mechanism, and the cohort comparison of Section VI tests that mechanism's timing directly.

III.D The other side of the market: redistribution, not destruction

The decomposition passes a sign-reversal test specified before the women's series was estimated. Carried out identically for young non-college educated women, equation (1) yields a positive between component of $+1.73$ to $+1.86$ points, because single non-college educated women participate 15 to 16 points above their married counterparts. Aggregate participation among young non-college educated women rose 3.21 points over the window

TABLE 2. Partnership composition and cell participation, young men and women by education and partnership, 2005–07 vs. 2015–19

	Share of all young (%)		Cell LFP (%)	
Cell	2005–07	2015–19	2005–07	2015–19
Men				
Non-college single (NCS)	55.80	59.68	84.90	82.04
Non-college married (NCM)	25.84	15.96	93.17	92.98
College single (CS)	11.69	18.01	95.46	95.72
College married (CM)	6.67	6.34	97.12	97.64
Women				
Non-college single (NCS)	43.28	46.80	78.29	79.30
Non-college married (NCM)	30.57	20.01	62.00	64.15
College single (CS)	14.24	21.57	94.90	95.81
College married (CM)	11.91	11.62	80.18	83.02

Notes. Shares of the within-gender population aged 21–30, not enrolled; ACS person weights. As a share of the whole young population the non-college married cell falls 9.9 points for men and 10.6 for women (within the non-college group alone, the married share falls 10.6 and 11.4 respectively); the college-married share is approximately unchanged for both, so the marriage decline is non-college-specific on both sides. Marriage–LFP gap among non-college educated men: +8.3 points in 2005–07, +10.9 in 2015–19. Source: authors' calculations from ACS via IPUMS.

TABLE 3. Within/between decomposition of the change in young non-college labor-force participation, 2005–07 to 2015–19 (percentage points)

	Men		Women	
	Laspeyres	Paasche	Laspeyres	Paasche
Aggregate change	−3.17		+3.21	
Within (cell rates)	−2.01	−2.29	+1.48	+1.35
Between (composition)	−1.15	−0.87	+1.73	+1.86
Between as share of total	36%	28%	54%	58%

Notes. Decomposition (1) of the aggregate participation change for non-college educated men and women aged 21–30. The *within* term holds the married/single composition fixed and moves cell participation rates; the *between* term holds rates fixed and moves composition. The Laspeyres reading weights the within term by base-period (2005–07) shares and the between term by end-period rates; the Paasche reading reverses both; each is algebraically exact. The between (composition) component is 28–36 percent of the men's 3.2-point decline and 54–58 percent of the women's 3.2-point rise. Source: authors' calculations from ACS via IPUMS.

in which the men's fell 3.17, and the between component accounts for 54 to 58 percent of that rise. The single share rose 10.6 points among non-college educated men and 11.4 among non-college educated women, a near-parallel movement of the two sides of one market. In measured participation, then, the digital-era marriage decline moved young non-college labor supply in opposite directions for the two genders: down for men, up for women.

The reversal is one of measurement, not of effort, and the time-use record shows it directly, in both the pre- and post-smartphone windows. Table 4 decomposes the average day of young men and women, by education and partnership status, into home production (household activities and the care of household members, which the labor-force statistics do not count) and market work, which they do, separately for 2005–2007 and 2015–2019. The married total-work surplus is positive in every cell of both windows. Among young non-college educated women in 2015–2019 the married supply 131 more minutes per day of home production than the single and 61 fewer of market work, for 70 more minutes of *total* work; among young non-college educated men the married supply more of both, 60 minutes of home production and 80 of market work, for 141 more in total. Both genders work more in total when married, before the smartphone as after. What reverses across the genders is not the quantity of work but the margin on which it falls: market hours for men, home production for women. Only the first is counted as participation. The married–single participation gap, positive for men and negative for women, is in this light a statement about where work is recorded, not how much is performed.

Splitting the windows shows this surplus did not shrink as mobile internet diffused; if anything it grew. Among young non-college educated adults the married–single total-work gap widened between the windows, from $+18$ to $+70$ minutes for women and from $+112$ to $+141$ for men. We read this descriptively rather than as a digital effect, for two reasons. It is a within-cell change in hours, a different object from the participation–composition channel the paper identifies: it concerns how much the married and the single each work, not the reallocation of the population across the partnership margin that moves the aggregate. And the two windows straddle the Great Recession (2005–2007 precedes it, 2015–2019 follows the recovery), so the widening blends a differential cyclical rebound of married and single hours with anything specific to the period; the levels show the gap growing as much through rising married hours as through any decline among the single. Consistent with that, the single non-college educated men whose market hours a leisure-luxury account would predict falling lose only a few minutes across the windows, within sampling noise and in line with the flat market-work path of Section IV. The durable statement is the level one: married young adults work more in total than single ones in both windows, so the participation reversal is a matter of accounting, not of effort.

The pattern is general: the total-work surplus of the married is positive in every age and education group we examine, including young college men and women, while its market-work component reverses sign by gender at working ages, both gaps largest among the young, where partnership formation and its labor consequences concentrate. This also disciplines a selection reading of the marriage–LFP gap. Were the gap pure selection, a fall in marriage would merely relabel men whose participation was already low; it would not move the aggregate, would not move the women's aggregate in the opposite direction, and certainly would not raise total work for both genders at once. A reorganization of labor supply within and across couples predicts the observed symmetry; a relabeling of fixed types does not. The model of Section V below takes this position exactly: the provision channel raises total productive work for both genders, and its measured market component can move in opposite directions across them because one margin is counted and the other is not. The claims here concern participation rates and time use, and make no statement about earnings or welfare.

IV. TIME ALLOCATION: EVIDENCE FROM THE ATUS

IV.A Solitary screens and the withdrawal from friends

The time-use record speaks to the mechanism by who is present, not merely to what activity is recorded. Table 5 reports the change in daily minutes of solitary digital leisure and of leisure with friends, between the pre- and post-smartphone windows, across all eight education-by-gender-by-partnership cells. The pattern is a loose-attachment gradient. Single non-college men increased solitary screen leisure by 34.3 minutes per day and single college men by 29.0, against five to six minutes for married men of either schooling and one to nine minutes for women. The comparison is the discriminating one: total digital

TABLE 4. Market and nonmarket work by gender, education, age, and partnership status, before vs. after the smartphone (minutes per day)

	Total work (level)				Married − single gap								
	Married		Single		Home			Market			Total		
	Bef.	Aft.	Bef.	Aft.	Bef.	Aft.	Δ	Bef.	Aft.	Δ	Bef.	Aft.	Δ
Panel A. Non-college (no bachelor's)													
Women													
All ages	370	357	309	295	+72	+65	−7	−11	−3	+7	+61	+62	+1
Ages 25–54	418	424	393	387	+73	+75	+2	−48	−37	+11	+26	+38	+12
Ages 21–30	406	444	388	374	+111	+131	+20	−94	−61	+33	+18	+70	+52
Men													
All ages	379	366	313	300	+30	+38	+8	+36	+28	−8	+66	+66	+0
Ages 25–54	451	451	371	363	+24	+37	+13	+55	+51	−4	+79	+88	+9
Ages 21–30	465	478	353	338	+45	+60	+16	+68	+80	+13	+112	+141	+28
Panel B. College (bachelor's or higher)													
Women													
All ages	423	398	377	360	+97	+82	−15	−50	−43	+7	+47	+38	−8
Ages 25–54	456	447	421	412	+121	+107	−14	−87	−72	+15	+34	+35	+1
Ages 21–30	430	438	410	386	+139	+130	−9	−119	−77	+41	+20	+52	+32
Men													
All ages	423	404	386	366	+42	+45	+3	−4	−7	−3	+38	+37	+0
Ages 25–54	481	480	427	405	+51	+58	+6	+2	+18	+16	+53	+75	+22
Ages 21–30	466	429	392	370	+47	+55	+8	+27	+4	−23	+74	+59	−15

Notes. ATUS 2003–2023 Activity Summary file; not enrolled; ATUS final weights. "Bef." is the 2005–2007 window, "Aft." the 2015–2019 window, and Δ the after-minus-before change. The two "Total work" columns are mean daily minutes for married and single respondents; the gap columns are the married-minus-single difference in home production, market work, and total work. Home production is household activities and care of household members (categories 02–03, including childcare); market work is the working category; total work is their sum. "Married" is spouse present; "single" is the complement; "All ages" is 15 and over. The pre-window here is 2005–2007, matching the ACS labor windows of Tables 1–3; the pure time-use decompositions of Section IV (Tables 5–6) use the fuller ATUS 2003–2007 window. The Total gap equals Married minus Single total work up to rounding; a positive Δ means the married–single total-work gap widened across the windows. College ages-21–30 cells rest on small samples and are noisier. Source: authors' calculations from ATUS public-use files.

TABLE 5. Change in solitary and non-solitary digital leisure and leisure with friends, by cell, pre- to post-smartphone windows (minutes per day)

	Non-college				College			
	Men		Women		Men		Women	
	Single	Married	Single	Married	Single	Married	Single	Married
Digital leisure								
Solitary	+34.3	+6.3	+7.1	+1.6	+29.0	+4.8	+8.9	+1.4
With others	+0.5	+2.4	+0.8	+4.2	+5.4	+3.0	−0.9	+6.3
Leisure with friends	−20.4	+2.8	−18.7	−2.4	−21.4	+0.9	−18.2	−4.1

Notes. ATUS 2003–2023, activity and who files; ages 21–30, not enrolled; ATUS final weights. Pre window 2003–07, post window 2015–19. Cells partition the sample by education (non-college: no four-year degree), gender, and partnership (married: spouse present; single: its complement). Solitary digital leisure is screen-leisure activity (t120307, t120308) with every who-record coded Alone (18/19); non-solitary digital leisure is the same activity with at least one other person present; leisure with friends is category-12 leisure with a Friends (54) record present. The solitary-screen surge is concentrated among *single* men (about +34 minutes for the non-college and +29 for the college) and is almost entirely solitary: of those increases only +0.5 and +5.4 minutes came with anyone else present. The withdrawal from friends falls on the single of both genders (−18 to −21 minutes) and spares the married. Source: authors' calculations from ATUS public-use files.

leisure rose for nearly every group as smartphones diffused, but solitary digital leisure, screen time with no other person present, surged for single men in particular. The surge was, moreover, almost entirely solitary: of the 34.3 extra daily minutes of screen leisure for single non-college educated men, all but half a minute came with no one else present, and for single college men 29.0 of 34.4. Screens gave these men not more time with others but more time alone: the footprint of a withdrawal from social time, not of a richer social diet. The withdrawal from friends, by contrast, is shared: single people of both genders lost eighteen to twenty-one minutes per day with friends, while the married were spared.

Two features organize the interpretation. First, the reallocation of social time is an order of magnitude larger than the reallocation of labor supply: these men moved roughly half an hour per day toward solitary screens and twenty minutes away from friends, while their participation rate over the same interval continued a pre-existing decline of less than two-tenths of a point per year. Whatever mobile internet did to this population, it did far more to its social life than to its hours of work, consistent with a marriage channel and not with a direct labor-supply channel. Second, the asymmetry across cells is itself evidence.

The withdrawal from friends is general to the single of both genders; the move into solitary screen leisure is specific to single *men*, consistent with the gaming-intensive composition of the activity category and with the male concentration of the labor-supply facts the paper studies.

IV.B Where the screen time came from

The asymmetry above shows that the additional screen time was solitary and came partly at the expense of friends; it does not show what else it displaced, and in particular whether it displaced work. If solitary screens drew their time from work, the leisure-luxury account of the participation decline would be confirmed in this cell directly. Table 6 accounts for the full day, for every education-by-gender-by-partnership cell, and shows they did not. Between the pre- and post-smartphone windows, single non-college educated men added 35 minutes per day of screen leisure, almost all of it solitary: 34 of the 35 minutes came with no one else present. Market work supplied essentially none of it: it fell about four minutes, against a daily total near 280 and well within sampling noise. The increment came instead from television, which fell twenty-six minutes, and from in-person socializing, down eleven; the day shifted further toward sleep, up twenty-two. The cross-cell pattern locates the shift: the rise in screen leisure is large for single men (35 minutes for the non-college, 34 for the college) and small for married men and for women of either schooling, roughly eight minutes or less, the same loose-attachment gradient the marriage channel predicts. Nowhere does a large screen gain coincide with a fall in work; among single non-college educated men in particular, the extra screen time grew at the expense of other leisure and the company of others, not the job.

This is the discriminating test, because the leisure-luxury mechanism operates within the person: a cheaper, more absorbing leisure technology draws a given individual's time from work toward the screen. In the cell where screen time rose most, it drew none. The market-work margin of single non-college educated men did not move as their screen time more than doubled, the minutes counterpart of the flat within-cell participation path of Section III, attachment to work unbroken in hours as in participation. The aggregate

TABLE 6. Where the time went: change in the daily time budget across cells, pre- to post-smartphone windows (minutes per day)

	Non-college				College			
	Men		Women		Men		Women	
	Single	Married	Single	Married	Single	Married	Single	Married
Screen leisure, solitary	+34	+6	+7	+2	+29	+5	+9	+1
Screen leisure, non-solitary	+1	+2	+1	+4	+5	+3	−1	+6
Market work	−4	+3	−19	−6	+2	−36	−14	0
Television	−26	−19	−9	−20	−3	−3	+7	+11
In-person socializing	−11	0	−10	−8	−6	+20	−12	−14
Phone calls	0	0	−1	0	−2	+1	−5	−1
Sleep & personal care	+22	+6	+26	+1	+18	+21	+40	+26
All other	−16	+1	+6	+28	−44	−11	−24	−30
n (2003–07)	1446	1130	1878	1794	486	412	584	772
n (2015–19)	959	467	1018	667	520	298	544	497

Notes. ATUS 2003–2023; respondents aged 21–30; ATUS final weights; the same sample and construction as Table 5. Cells partition the sample by education (non-college: no four-year degree), gender, and partnership status (married: spouse present in the household; single: no spouse present, the closest Activity-Summary proxy for the legal-marital definition used elsewhere in the paper). Each entry is the change in mean daily minutes between the 2003–07 and 2015–19 windows; within a cell the eight activity rows partition the 1,440-minute day, so each column nets to approximately zero (up to rounding). Screen leisure is games (t120307) and computer use for leisure (t120308), split by the ATUS "who" file into solitary minutes (every co-present record coded alone, codes 18–19) and non-solitary minutes (at least one other person present), exactly as in Table 5; the remaining rows are the corresponding daily totals for market work, television, in-person socializing, phone calls, sleep and personal care, and an all-other residual. College cells, and married cells generally, are small in the post window (*n* as low as 298), so individual entries there carry substantial sampling noise; the robust feature is the solitary screen-leisure row. Source: authors' calculations from ATUS public-use files.

decline in young non-college educated men's work hours that the leisure-luxury literature documents is not thereby denied, but it is relocated: because married men work substantially more than single men, the fall in the married share lowers average hours across the group through the same composition that lowers participation, with each cell's hours little changed. The reallocation of social time this paper places at the center of the marriage decline is visible in the budget; a reallocation from work to screens, in the cell where screens grew, is not.

V. A MODEL OF RELATIONAL CAPITAL AND LABOR SUPPLY

The facts of the preceding sections rest on two regularities this section explains from a single primitive: married men participate in the labor force more than single men, and the married share of young men fell sharply over the smartphone era. We build on the household-production model of Hudson and Moscoso Boedo (2026b), in which the falling relative price of digitally mediated connection lowers households' investment in relational capital; that model takes labor supply as fixed. We endogenize it, and show that the same gross complementarity between the digital bundle and time that drives the partnership and fertility results there generates a *positive* relationship between relational capital and market work here. The model rationalizes the married–single participation gap documented in Section III and isolates the channel through which a depth-reducing technology shock lowers aggregate participation, which is composition: the shock moves the population across a fixed participation gradient rather than bending the gradient itself.

V.A The partnership channel

Durable partnerships rest on a stock of relational capital: shared history, mutual knowledge, trust, the accumulated investment of sustained in-person time. That stock is built primarily during early adulthood and primarily through direct, co-present interaction, and forming a marriage or marriage-like partnership grows more likely the more of it a person has accumulated. Mobile internet lowers the price, and raises the immediate return, of a competing use of social time: broad, shallow, digitally mediated connection. Time that would once have gone to the few deep, co-present relationships that build partnership-grade relational capital is reallocated toward many low-investment digital ties. The substitution is asymmetric: the digital good substitutes well for stimulation, validation, and ambient companionship, but poorly for the embodied, high-investment interaction that builds the depth a partnership requires. Relational capital therefore accumulates more slowly for the cohorts whose formative years coincide with cheap, ubiquitous mobile internet, fewer of them build the depth a partnership requires, and the cohort's marriage rate falls.

The mechanism has three empirical signatures, two of them already visible in the facts above and the third the subject of the cohort design below. The first concerns the margin: the effect operates on whether a marriage forms, not on the labor supply of men who remain single, implying a fall in the married share and through it a compositional decline in participation, with no necessary break in single men's own participation rate, the joint pattern Section III documents. The second concerns behavior: because the mechanism runs through the reallocation of social time, it leaves a footprint in time use (more time alone with a screen, less with friends), which the diaries of Section IV show. The third concerns timing: the effect should concentrate in the cohorts whose formative years overlapped the diffusion of mobile broadband, the variation the cohort comparison of Section VI exploits.[1]

V.B Environment

The economy is a unit continuum of agents indexed by i, identical in every primitive but one: the weight each places on the fertility value of relationships. Agent i draws this weight γ_i from a lognormal distribution F,

$$\log \gamma_i \sim \mathscr{N}\big(\log \tilde{\gamma},\ \sigma_\gamma^2\big), \tag{2}$$

with median $\tilde{\gamma}$ and log-dispersion σ_γ, both taken from Hudson and Moscoso Boedo (2026b); $\tilde{\gamma}$ is identified there from children per adult and σ_γ from the dispersion of close ties in the General Social Survey. This is the model's one source of heterogeneity; it is what gives the model a distribution of relationship depth, and through it a married share. We carry the index i through the agent's problem and aggregate over F at the end.

Each period agent i divides a unit of time among investment in relationships h_i, market work ℓ_i, and outside-option time x_i,

$$h_i + \ell_i + x_i = 1. \tag{3}$$

[1]The model is stationary and carries no age structure, so it does not by itself predict that the response concentrates in the exposed cohorts; that concentration is empirical content the cohort design supplies, and it reflects what the aggregate model abstracts from: relational capital is built primarily in adolescence and early adulthood. The two exercises are complementary: the calibration asks whether the aggregate magnitudes line up, the cohort comparison whether the response falls where formative exposure predicts.

In that model, ℓ_i is fixed; here it is a choice, and (3) is the only constraint. Relational capital D_i is a stock built only by direct co-present time and subject to depreciation $\delta \in (0,1)$,

$$D_i' = (1-\delta)D_i + h_i, \qquad \text{so that in stationary state} \qquad h_i = \delta D_i. \tag{4}$$

Fertility within a relationship is the saturating function $b(D) = 1 - e^{-D/\kappa}$, and a relationship of depth D supports a household of size

$$n(D) = 1 + b(D), \tag{5}$$

the agent together with children in expected number $b(D)$. The functions b, n, and π below are common to all agents; agents differ only in the depth D_i at which they are evaluated. Household size is the single object the model adds to the structure of Hudson and Moscoso Boedo (2026b), and it is built from the existing fertility primitive: deeper relationships sustain larger households. Household size grows smoothly with depth. Partnership is likewise smooth in depth rather than a threshold event: a relationship of depth D is realized as a partnership with probability $\pi(D)$, a logistic increasing in depth,

$$\pi(D) \;=\; \big[1 + \exp\!\big(-(D - D_0)/s\big)\big]^{-1}, \tag{6}$$

with location D_0 and scale s calibrated to the married share in Section V.F. The married share is the average of $\pi(D_i)$ across agents, made precise in (21); the smoothness is substantive, since a knife-edge threshold in depth would force that average to respond far too sharply to a common downward shift in the depth distribution, overshooting the observed decline, whereas the gradual logistic matches it (Section V.F).

Market income $w\ell_i$ finances the digital quality bundle, which must provide for a household of $n(D_i)$ members; per-capita quality is income divided among the household,

$$\widetilde{Q}_i = \frac{w\ell_i}{p}\,\psi(A)\,\frac{1}{n(D_i)} \;\equiv\; \frac{\xi\,\ell_i}{n(D_i)}, \qquad \xi \equiv \frac{w}{p}\,\psi(A), \tag{7}$$

where p is the hedonic-quality-adjusted price of the bundle, $\psi(A)$ the network multiplier in peer engagement A, and ξ collects the three exactly as in the baseline model and is common across agents. Larger households dilute the per-capita bundle; this is the provision motive, and it is the entire content of the extension. Per-capita consumption services aggregate the bundle and the agent's own leisure time through the same constant-elasticity technology, with elasticity of substitution $\sigma \in (0,1)$ and quality weight α,

$$c_i = \left[\alpha\, \widetilde{Q}_i^{\rho} + (1-\alpha)\, x_i^{\rho}\right]^{1/\rho}, \qquad \rho \equiv \frac{\sigma - 1}{\sigma} < 0, \tag{8}$$

in which $\sigma < 1$ makes the bundle and time gross complements. Flow utility is separable in consumption and the fertility value of depth,

$$U_i = \log c_i + \gamma_i \log b(D_i), \tag{9}$$

the only place the type γ_i enters. Because (9) and the depth technology are unchanged, the quantitative discipline that model places on the fertility and partnership margins carries over intact; the present extension alters only what the budget must finance, and only the labor margin.

V.C The household's problem

Agent i enters the period with relational capital $D_{i,-1}$ and chooses (h_i, ℓ_i, x_i) subject to the time constraint (3). Because household size reflects relationships already formed, the per-capita budget (7) is evaluated at the predetermined size $n(D_{i,-1})$; current investment h_i adds to the stock that supports the household in the following period. Writing $V(D_{i,-1};\gamma_i)$ for the value of a type-γ_i agent entering with stock $D_{i,-1}$ and $\beta \in (0,1)$ for the discount factor, the agent's problem is the Bellman equation

$$V(D_{i,-1};\gamma_i) \;=\; \max_{h_i+\ell_i+x_i=1} \left\{\log c_i + \gamma_i \log b(D_i)\right\} \;+\; \beta\, V(D_i;\gamma_i), \quad D_i = (1-\delta)D_{i,-1} + h_i, \tag{10}$$

with c_i given by (8) evaluated at $\widetilde{Q}_i = \xi \ell_i / n(D_{i,-1})$. Letting μ_i be the multiplier on the time constraint, the first-order conditions for market work and leisure equate their marginal contributions to consumption,

$$\frac{1-\omega_i}{\ell_i} \;=\; \mu_i, \qquad\qquad \frac{\omega_i}{x_i} \;=\; \mu_i, \tag{11}$$

where ω_i and $1-\omega_i$ are the leisure and bundle shares of the consumption aggregate,

$$\omega_i = \frac{(1-\alpha)\,x_i^{\rho}}{\alpha\,\widetilde{Q}_i^{\rho} + (1-\alpha)\,x_i^{\rho}}, \qquad 1-\omega_i = \frac{\alpha\,\widetilde{Q}_i^{\rho}}{\alpha\,\widetilde{Q}_i^{\rho} + (1-\alpha)\,x_i^{\rho}}, \tag{12}$$

and the condition for relationship investment equates the shadow value of time to the marginal value of depth,

$$\mu_i \;=\; \gamma_i\,\frac{b'(D_i)}{b(D_i)} \;+\; \beta\,V'(D_i;\gamma_i). \tag{13}$$

The marginal value of relational capital follows from the envelope condition,

$$V'(D_{i,-1};\gamma_i) \;=\; (1-\delta)\,\mu_i \;-\; (1-\omega_i)\,\frac{b'(D_{i,-1})}{n(D_{i,-1})}, \tag{14}$$

whose first term is the depreciation-adjusted shadow value of the stock, exactly as in the baseline, and whose second term is the one feature the budget extension adds: a deeper relationship entering the period supports a larger household, diluting per-capita consumption, so accumulated depth carries a provision cost. Substituting (14) into (13) and imposing stationarity gives the depth Euler equation

$$\mu_i\,\big[1-\beta(1-\delta)\big] \;=\; b'(D_i)\Big[\frac{\gamma_i}{b(D_i)} \;-\; \beta\,\frac{1-\omega_i}{n(D_i)}\Big], \tag{15}$$

which pins down the stationary depth of a type-γ_i agent; writing that solution as $D(\gamma_i)$, the policy $\gamma \mapsto D(\gamma)$ together with F delivers the cross-sectional distribution of depth that the composition channel acts on. The work and leisure conditions (11) are the new content of this section, and they fix the labor margin independently of the depth equation: because consumption depends on h_i only through the predetermined stock, the split of the remaining time between work and leisure analyzed in the next subsection is governed by (11) alone and is unaffected by the provision cost in (15).

V.D Relational capital and labor supply

We characterize the effect of relational capital on market work in two steps: the labor–leisure margin holding the time budget fixed, which isolates the provision effect, and the full effect once the time cost of building depth is admitted.

Fix the time available to agent i for work and leisure, $T_i \equiv \ell_i + x_i = 1 - h_i$, and consider how the agent allocates it. Both ℓ_i and x_i enter utility only through consumption, so the allocation equates their marginal contributions, $\partial c_i/\partial \ell_i = \partial c_i/\partial x_i$. Substituting (7) into (8) and simplifying, the work-to-leisure ratio is

$$\frac{\ell_i}{x_i} = \left(\frac{\alpha}{1-\alpha}\right)^{\sigma} \left(\frac{n(D_i)}{\xi}\right)^{1-\sigma}, \tag{16}$$

so that the elasticity of the work-to-leisure ratio with respect to household size is exactly

$$\frac{\partial \log(\ell_i/x_i)}{\partial \log n} = 1 - \sigma, \tag{17}$$

the same for every agent. The sign of this elasticity is governed entirely by gross complementarity. Under $\sigma < 1$ it is positive: a larger household dilutes the per-capita quality bundle, and because the bundle and time are complements, the scarcer bundle raises the marginal value of the income that restores it, drawing time toward market work. At the Cobb–Douglas knife-edge $\sigma = 1$ household size drops out of the labor allocation; under gross substitutes $\sigma > 1$ the sign reverses. Since $n(D)$ is increasing in D, (17) states that, holding the time budget fixed, greater relational capital raises market work, and it does so through the same elasticity that governs the partnership and fertility results of the baseline model.

Proposition 1 (The provision channel)**.** *Holding the time available for work and leisure fixed, the work-to-leisure ratio satisfies* (16)*, with elasticity* $1-\sigma$ *with respect to household size* $n(D)$*. Under gross complements* $(\sigma < 1)$ *the ratio is increasing in* $n(D)$*, and hence in* D*: greater relational capital raises market work. The effect vanishes at* $\sigma = 1$ *and reverses for* $\sigma > 1$.

The provision channel is offset by a time cost. Maintaining a deeper relationship requires more investment, $h_i = \delta D_i$ in stationary state, leaving less time for work and leisure, $T(D_i) = 1 - \delta D_i$. Writing $r_i \equiv \ell_i / x_i$ from (16) and $\ell_i = T(D_i)\, r_i / (1+r_i)$, the total effect of relational capital on agent i's market work is

$$\frac{d\ell_i}{dD_i} \;=\; \underbrace{\frac{T(D_i)}{(1+r_i)^2}\,(1-\sigma)\,\frac{b'(D_i)}{n(D_i)}\,r_i}_{\text{provision }(+)} - \underbrace{\delta\,\frac{r_i}{1+r_i}}_{\text{time cost }(-)}\;. \tag{18}$$

A model without the household in the budget is the case $n(D) \equiv 1$: the provision term vanishes, only the time cost survives, and relational capital unambiguously lowers market work. Such a model makes relationships pure competitors for time and predicts that single agents work more than partnered ones, the opposite of the gap in the data. With provision present, the net effect is positive whenever

$$\frac{(1-\sigma)\,b'(D_i)\,T(D_i)}{n(D_i)\,(1+r_i)} \;>\; \delta, \tag{19}$$

that is, when the complementarity is strong, the household grows steeply with depth, and depth is inexpensive to maintain. The married–single participation gap of eight to eleven points documented in Section III is the evidence that this inequality holds: in the relevant region provision dominates the time cost, and the magnitude of the gap disciplines how steeply household size rises with depth.

Put differently, endogenizing labor without the provision channel would make relational capital a pure time cost, so that agents with more relational capital work less. Once the household a deeper relationship supports must be provided for, the provision channel of Proposition 1 reverses the sign under condition (19), which the observed participation gap implies: agents with more relational capital work more.

The interpretation of the gap that (18) supplies is causal (partnership raises the marginal value of market income through the household it supports), but the composition accounting of the empirical sections does not require committing to it. Whether the measured gap reflects this channel or selection, the change in the married share moves aggregate participation through the same arithmetic; the model supplies one mechanism consistent with that arithmetic and ties it to the relational-capital primitive.

V.E The technology shock and the composition channel

The technology shock acts as in the fertility model and reaches labor supply through depth. A fall in the relative price of the digital bundle, or a rise in its network value, raises ξ and with it every agent's effective bundle $\widetilde{Q}_i$; under gross complements this raises the leisure share ω_i and lowers the opportunity cost of leisure relative to depth, so the depth D_i falls. Lower depth contracts the household, $n(D_i)$ declines, and by Proposition 1 the provision pull on market work weakens. The chain, holding for every agent, is

$$\xi \uparrow \Longrightarrow D_i \downarrow \Longrightarrow n(D_i) \downarrow \Longrightarrow \ell_i \downarrow, \tag{20}$$

under condition (19). The reallocation of social time away from depth-building investment that lowers fertility there lowers market work here, through the household that depth supports: a single mechanism on two margins of the same household-production problem.

Two features of (20) connect the model to the empirical record. First, the positive gradient of work in depth, $d\ell_i/dD_i > 0$, is precisely what transmits the depth decline to labor: an economy in which relational capital is a pure time cost, with $d\ell_i/dD_i < 0$, would respond to a fall in depth with a *rise* in market work. The provision channel is therefore not in tension with the participation decline but is its mechanism, and it is the structural counterpart of the married–single gap. Second, the shock lowers depth for every agent ($D(\gamma_i)$ falls for all i in (20)), so the entire cross-sectional distribution of depth shifts down, and with it the married share. Writing the married share as the average of the partnership probability (6) across agents,

$$m \;=\; \mathbb{E}_i\big[\pi\big(D(\gamma_i)\big)\big] \;=\; \int \pi\big(D(\gamma)\big)\,dF(\gamma), \tag{21}$$

a fall in every $D(\gamma_i)$ lowers m, because π is increasing. The participation consequence is then composition. Let p_M and p_S be the participation rates of partnered and single agents; aggregate participation and its response to the shock are

$$L \;=\; m\,p_M + (1-m)\,p_S, \qquad dL \;=\; (p_M - p_S)\,dm, \tag{22}$$

so the shock lowers participation in proportion to the married-share decline it produces, scaled by the partnered–single gap, with the within-cell rates p_M and p_S held fixed. This is the central empirical finding in the model's own terms: participation moves through the population's redistribution across the partnership margin, not through any change in the gradient $p_M - p_S$ itself. The provision channel of Proposition 1 is what rationalizes $p_M > p_S$; the quantitative conversion in (22) nonetheless uses the measured gap, not the model's, so the participation result does not depend on the model reproducing the gap's size.

The model is written for an agent whose productive time is measured by market work, which is the case for men. For women, productive time includes home production that the labor-force measure does not count; the provision channel raises total productive work for both genders, while its measured market component can move in opposite directions across the genders because one margin is counted and the other is not. The model takes no stand on the allocation of work within the household: it requires only that the household a deeper relationship supports raise total productive work, which market participation captures cleanly for men. The quantitative analysis is therefore conducted on men, for whom the object the model describes and the object the data record coincide. Section III confirms the premise directly in time-diary data: married men and women both supply more total work than their single counterparts, the additional work falling on market hours for men and on uncounted home production for women.

V.F Calibration: marriage and participation as the price falls

The facts route the participation decline through the married share, and the model of Section V ties the married share to the falling price of the digital bundle. This section asks the quantitative question that pairing raises: driven by the same falling smartphone hedonic price that explains fertility in the baseline, does the model reproduce the magnitudes the facts establish, the married-share and participation declines and their timing, with nothing fit to labor?

The model's structural parameters are taken from Hudson and Moscoso Boedo (2026b), but the discipline behind them is empirical, and Panel A of Table 7 lists the moment that identifies each: the elasticity σ from the response of the outside-option time share to the price, the quality weight α from its level, the depreciation rate δ from the half-life of relational capital, the birth curvature κ from the fertility–depth gradient, the dispersion σ_γ from the distribution of close ties in the General Social Survey, and the network multiplier from the response of peer engagement. Two objects are set here, to moments of our own sample: the fertility constant μ_0 to the 2007 outside-option time share, and the partnership logistic (6), whose location D_0 and scale s map the model's depth distribution into a married share through (21) and are calibrated to the married share at the start of the transition, 2007, and at its observed terminal level, the 2015–19 window value of 21.1 percent, which the observed share essentially holds thereafter (20.0–20.4 percent through 2023) and which the transition reaches by its 2024 terminal date. The labor extension of Section V adds no parameter; it changes the depth Euler (15) only through the small provision term. For labor the exercise is out of sample: a model disciplined by time use, fertility, and marriage, asked to deliver participation.

We solve the model dynamically. Rather than compare two steady states, we compute the economy's perfect-foresight transition as the wage-deflated, quality-adjusted price of a smartphone falls year by year, raising the effective digital bundle $\widetilde{Q}$ along the path in Panel A of Figure 2, anchored to the published 2007 and 2024 bundle values, with the

TABLE 7. Calibration: parameters, identifying moments, and the dynamic response

Panel A. Parameters and identifying moments

Parameter	Symbol	Value	Identifying moment
Elasticity of substitution	σ	0.53	share response to price
Quality weight	α	0.64	outside-option share, level
Depth depreciation	δ	0.15	relational-capital half-life
Birth curvature	κ	10	fertility–depth gradient
Discount factor	β	0.96	standard
Relational-capital weight, median	$\tilde{\gamma}$	1.32	children per adult, 2007
Relational-capital dispersion	σ_γ	0.55	GSS close-tie distribution
Network multiplier	$\psi_{\max}, k, A^*$	$0.71, 20, 0.72$	peer-engagement response
Fertility constant	μ_0	0.043	outside-option share, 2007
Partnership mapping	D_0, s	4.8, 1.3	married share, 2007 & 2015–19

Panel B. Dynamic transition, 2007–2024

	Data	Model
Calibration targets		
Outside-option time share, 2007	0.49	0.49
Married share, 2007	31.6%	31.6%
Married share, 2015–19	21.1%	21.1%
Untargeted outcomes		
Conditional fertility, 2007–24	−18.5%	−17.9%
Participation, implied by composition	−0.9 to −1.2 pp	−1.0 pp

Notes. Panel A: σ through the network multiplier are taken from Hudson and Moscoso Boedo (2026b), identified there by the listed moments; μ_0 and the partnership mapping are calibrated here to the listed moments of our sample. The partnership mapping is a logistic in relationship depth, $p(D) = [1 + \exp(-(D - D_0)/s)]^{-1}$, set to the observed married share in 2007 and in the 2015–19 window (21.1 percent; the observed share is essentially flat thereafter, 20.0–20.4 percent through 2023), which the transition reaches by its 2024 terminal date. The labor margin adds no parameter. Panel B: the model is solved as a perfect-foresight transition as the effective digital bundle $\widetilde{Q}$ rises from 3.3 (2007) to 90 (2024) along the path of Figure 2. The three targets are matched by construction. Conditional fertility is untargeted (US children per adult aged 25–44). Participation is the composition channel of equation (1), the married-share decline times the period-average married–single gap, which delivers the between-component, about a third of the −3.2-point total; the within-cell remainder lies outside the channel. Married share and participation are for non-college educated men aged 21–30. Source: authors' calculations; structural parameters from Hudson and Moscoso Boedo (2026b).

front-loaded profile of the device-price decline in between. Relationship capital is a stock that depreciates at rate δ and cannot be shed faster than it depreciates, so when the price falls the depth distribution slides down only gradually, lagging the price (Panel A). That lag is the dynamic content a steady-state comparison misses: the price has largely finished falling by the mid-2010s, but the margins keep moving into the 2020s as relational capital

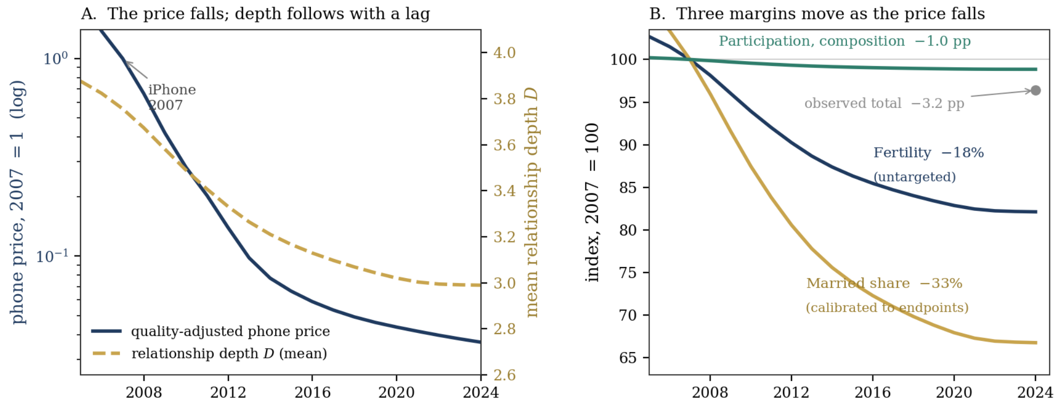

FIGURE 2. The smartphone-era price shock as a dynamic transition. Panel A: as the wage-deflated, quality-adjusted price of a smartphone falls (navy, log scale, 2007 = 1), the stock of relational capital depreciates toward its new level only gradually (gold), lagging the price: the price has largely finished falling by the mid-2010s while depth keeps declining into the 2020s. Panel B: the same falling price moves three margins over time, indexed to 2007 = 100: conditional fertility (untargeted by the calibration), the married share (its mapping from depth calibrated to the 2007 level and the 2015–19 window level), and, through the composition channel of equation (1), labor-force participation. The participation decline the channel delivers, about one point, is the between-component; the observed total (−3.2 points, marker) includes the within-cell slide of the single. Conditional fertility falls 17.9 percent against an observed 18.5. The price path is anchored to the published 2007 and 2024 bundle values, with the device-price profile in between.

depreciates toward its new level.

The transition delivers three trajectories, in Panel B of Figure 2 and Panel B of Table 7. Conditional fertility falls 17.9 percent between 2007 and 2024, against an observed 18.5, and fertility is not a target. It falls out of the depth dynamics: because depth in 2007 is still elevated from the pre-iPhone era and declines through the entire window, the transition reproduces the *total* fertility decline, not the post-2007 marginal piece that a frozen-2007 steady state isolates. The married share, its mapping from depth calibrated to the 2007 level and the 2015–19 window level, traces a smooth decline whose *timing* is the model's: gradual, tracking the depreciation of relational capital rather than the faster fall in the price, and largely complete by the early 2020s. Participation follows by composition, equation (1): the married-share decline times the participation gap of Section III yields a fall of about 1.0 point, the between-component, roughly a third of the observed 3.2-point aggregate decline. The remainder is the within-cell slide of the single, which predates the smartphone and lies outside the digital channel. One falling price; three margins moving together, with the lag that a stock of relational capital imposes.

Three caveats bound the claim. First, the intermediate shape of the price path is a reconstruction; only its 2007 and 2024 endpoints are pinned to published values. The fertility trajectory's *magnitude* is robust to that shape, because relational capital adjusts most of the way to its 2024 level by 2024 whatever path the price takes, but the precise year-by-year timing in Figure 2 inherits the reconstructed profile. Second, the marriage magnitude is not a free prediction. The decline a given depth movement produces depends on the smoothness of the partnership mapping, which a sharp threshold would amplify far past the data; we therefore calibrate that mapping to the observed married-share endpoints, so the model's marriage content is its *direction and timing*, not its level. Fertility, by contrast, is untargeted and matches in magnitude. Third, the model's cross-sectional reach is narrower than its time-series reach: the provision channel of Proposition 1 predicts a married–single labor gap of the right sign but, at this calibration, only a few points on the intensive margin, against the larger participation gap in the data. The composition channel uses the empirical gap, not the model's, so the participation result does not rest on the model reproducing that gap. The model is asked to deliver the marriage decline's direction and horizon, and it does so only in part: the observed share falls fastest before the mid-2010s and is essentially flat thereafter, while the stock-adjustment logic of relational capital spreads the model's decline more evenly into the early 2020s, so the actual path is more front-loaded than the model's. The gap that converts marriage into participation is measured, not assumed.

VI. COHORT EXPOSURE AND THE MARRIAGE DECLINE

The calibration took the digital origin of the marriage decline as given, inheriting it from the fertility evidence. This section subjects that origin to a direct test. The decomposition isolated the post-2007 fall in the married share as the phenomenon driving the aggregate decline; the question here is whether mobile internet drove that fall. Because mobile coverage is correlated with population density, and dense and sparse places differ in family formation for reasons unrelated to technology, a comparison of high- to low-coverage *places* would load that difference onto the treatment. We obtain variation that does not track density from a within-place comparison of *cohorts*.

VI.A The comparison, in words

The thought experiment is simple. Two men grow up in the same commuting zone but are born a few years apart. Because 4G reached their zone on a fixed calendar date, the younger man spent more of his formative years (ages 15–21, when social habits and the capacity for partnership take shape) under local coverage than the older man did. We ask whether the younger, more-exposed cohort married at a lower rate than the older one, *within the same place*. Comparing cohorts inside a zone, rather than high- to low-coverage zones, is what removes the urban–rural gradient: a zone's industrial base, its marriage culture, and how early it was covered are all differenced out, because they are properties of the place and not of the cohorts being compared. What is left to drive a difference between cohorts of the same zone is that the rollout overlapped some birth cohorts' formative windows and not others'.

VI.B Data, sample, and construction

Marriage outcomes are from the pooled American Community Survey, 2005–2023, via IPUMS. We assign each respondent to a 1990 commuting zone (Tolbert and Sizer, 1996), a cluster of counties forming a local labor market whose geography is stable across the public-use boundary revisions of 2012 and 2022, through their PUMA, using the crosswalk vintage appropriate to the survey year, and we define their birth cohort as survey year minus age. A *cell* is a commuting zone $\times$ birth cohort: the ever-married rate of all men (women) aged 25–34 of that cohort observed in that zone, together with its sampling standard error. The cohort design departs from the labor analysis in two ways, both by necessity: it uses ages 25–34, since marriage is rare and uninformative before the mid-twenties, and its baseline pools all education for power, with Section VI.F splitting the estimate by education and showing it concentrated among non-college educated men, the population of the labor analysis. The estimating sample is the 14 cohorts born 1983–1996 in the 493 commuting zones with at least 25 sampled men per cell: 6,436 cells, covering 2.10 million men, the median cell holding 123 respondents (the women's sample is 6,360 cells and 2.05 million respondents). Summary statistics are in Panel B of Table 1.

Local 4G coverage $\mathrm{Cov}_{z,t} \in [0,1]$ is the share of commuting zone z covered by fourth-generation mobile service in calendar year t, aggregated to commuting zones from the FCC's mobile-deployment filings, 2010–2019. Commercial LTE did not begin in the United States until December 2010, so coverage is a true zero before then. For a man in birth cohort c and commuting zone z, formative exposure is the mean local coverage over the formative years, ages 15–21, that fall within the coverage panel,

$$E_{zc} \; = \; \frac{1}{|W_{zc}|} \sum_{a \in W_{zc}} \mathrm{Cov}_{z,c+a}, \qquad W_{zc} = \{\, a \in \{15,\ldots,21\} : 2010 \leq c+a \leq 2019 \,\}, \tag{23}$$

with $E_{zc} \equiv 0$ when $W_{zc} = \varnothing$ and $c+a$ the calendar year cohort c reaches age a. Because commercial LTE did not exist before December 2010, every cohort whose formative years predate the rollout has $E_{zc} = 0$ exactly, a true zero rather than a missing value, and exposure then rises across successively younger cohorts at a pace fixed by how quickly each zone was covered. By construction, then, the variable that turns "on" across cohorts is the interaction of *where* (how rugged or remote the zone, hence how late its towers) with *when* (whether the rollout fell inside the cohort's window).

VI.C Specification

We read the cohort-by-cohort marriage response as an event study on commuting-zone $\times$ birth-cohort cells,

$$M_{zc} \; = \; \alpha_z + \delta_c + \sum_{k \neq k_0} \beta_k \left(\overline{\mathrm{Cov}}_z \cdot \mathbf{1}\{c = k\} \right) + \varepsilon_{zc}, \tag{24}$$

where M_{zc} is the ever-married rate of cohort c in zone z in the pooled ACS, $\overline{\text{Cov}}_z$ is the zone's 2013 4G coverage (a single cross-sectional index of how early and extensively the rollout reached zone z), α_z and δ_c are commuting-zone and birth-cohort fixed effects, and k_0 is an omitted pre-rollout reference cohort. Cells are weighted by the inverse of their squared sampling standard error, and standard errors are clustered by commuting zone (493 clusters for men, 492 for women). The two sets of fixed effects carry the identification: the zone effects α_z absorb everything fixed about a place, the density gradient included, and the cohort effects δ_c absorb national trends in marriage timing. What is left to move β_k is only the within-zone, across-cohort interaction of (23): how a zone's 2013 coverage maps into the marriage rate of cohort k relative to the reference. Equation (24) is the nonparametric counterpart of the exposure idea in (23): formative exposure rises with both the zone's coverage and the share of the formative window that postdates the rollout, so under the exposure logic the loading on $\overline{\text{Cov}}_z$ should be zero for the pre-rollout cohorts and rise across the exposed ones. Rather than impose that profile through a single E_{zc} regressor, we interact $\overline{\text{Cov}}_z$ with a full set of cohort indicators and let each cohort's loading β_k be read from the data, so the predicted flat-then-rising shape is tested rather than assumed.

The path $\{\beta_k\}$ carries a prediction that is also its own assumption check. For cohorts whose window closed before the rollout, high- and low-coverage zones are not yet distinguished ($E_{zc} = 0$ in both), so β_k should be zero; these pre-rollout coefficients are a parallel-trends test internal to the same regression. At the cohorts whose window caught the rollout, β_k should turn negative. The decisive feature is the shape. A smooth secular decline in marriage, the leading non-digital alternative and the one no cross-place comparison can exclude, tilts the whole cohort path; it cannot produce coefficients that are flat across the early cohorts and break at one particular birth year, whereas a force tied to the timing of the rollout does exactly that. The estimand is the effect of formative mobile-internet exposure on the marriage rate of the exposed cohorts, relative to cohorts of the same zones whose windows preceded it.

VI.D The marriage path: flat before, breaking at the exposed cohorts

Figure 3 reports the estimated path $\{\beta_k\}$ for the ever-married rate, and it has the predicted shape. Across the pre-rollout cohorts, those born 1983 through 1988, whose formative years predate commercial LTE, the coefficients are flat: the joint test of $\beta_k = 0$ for the 1983–1987 cohorts (measured against the 1988 reference) cannot reject ($F = 0.49$, $p = 0.78$ for men; $F = 0.48$, $p = 0.79$ for women), so the internal parallel-trends check passes. The 1989 and 1990 cohorts, whose formative years only grazed the rollout, sit near zero and insignificant. The 1989 cohort's window touches the panel only at age 21 in 2010, before commercial LTE existed, so its measured exposure is exactly zero; the 1990 cohort, whose window reaches 2011, carries mean exposure of about ten percent. Slight drift at the 1990 cohort is therefore the beginning of the dose response, not a pre-trend, and the formal pre-trend test is defined on the strictly unexposed cohorts. The path then breaks negative from the 1991 cohort onward, as formative-window overlap grows, reaching -0.066 for women and -0.040 to -0.043 for men at the most-exposed cohorts, with the exposed cohorts (1991–1996) jointly rejected at $p < 0.0001$ for both genders. The flat-then-break shape is the signature a smooth secular decline cannot reproduce. The male and female paths are parallel, both flat before the rollout and breaking at the exposed cohorts as two sides of one marriage market should, with the female decline somewhat larger in magnitude and consistent with the independent estimate on national data in Hudson and Moscoso Boedo (2026b).

Because exposure rises continuously across cohorts rather than switching on, no single break year need be sharp in every panel; the exact cohort at which each path turns reflects the group's exposure profile and sampling noise, and the identifying statement is the joint contrast between the strictly unexposed cohorts and the exposed set, not the location of any one kink. The exposure profile itself makes this concrete: mean formative exposure is exactly zero through the 1989 cohort, about ten percent for the 1990 cohort, and rises steeply from 0.13 to 0.55 across the 1991–1996 cohorts.

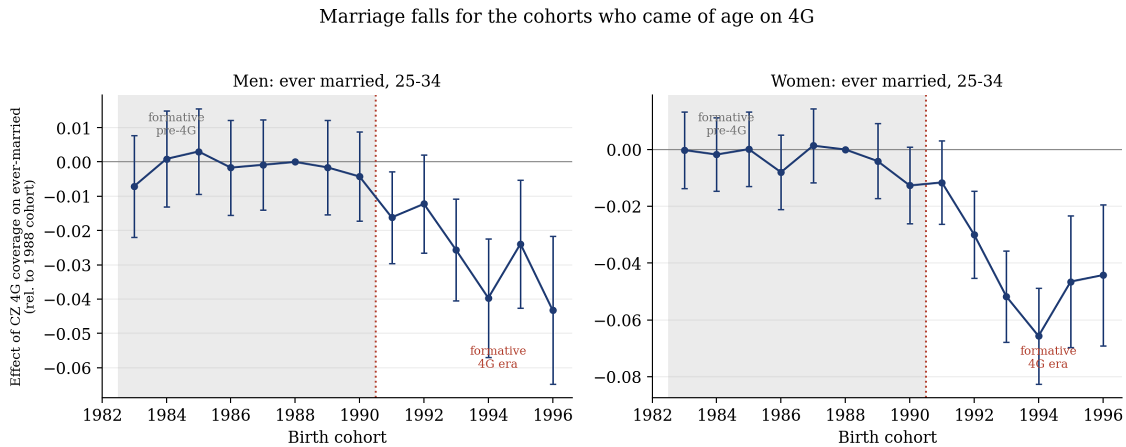

FIGURE 3. Cohort event study of 4G exposure on the ever-married rate (all men and women aged 25–34, all education levels). Left: men; right: women. The shaded region marks the pre-rollout cohorts, whose formative years (ages 15–21) largely predate commercial LTE; the coefficient path is flat across them and breaks sharply negative at the cohorts whose formative window caught the rollout. Bands are 95 percent confidence intervals.

VI.E Robustness

The break survives the confounds the design most naturally invites; Table 8 collects the checks. First, the flat pre-rollout path is a formal result, not an eyeballed one: the joint test of the pre-rollout cohort coefficients fails to reject zero for both genders, while the exposed-cohort coefficients are jointly significant at $p < 0.0001$. The flatness survives two further probes. Re-estimating equation (24) with the 1985 cohort rather than 1988 as the omitted reference leaves the most-exposed coefficients essentially unchanged (-0.043 for men, $t = -4.8$; -0.066 for women, $t = -6.9$) and the pre-rollout path flat (joint $F = 0.62$, $p = 0.65$ for men; $F = 0.56$, $p = 0.69$ for women), so the flat-then-break shape is not an artifact of the omitted category. And the sharper worry, that a break might in fact predate the rollout, has a direct test: restricting the sample to the pre-rollout and grazing cohorts alone (1983–1990) and replacing the cohort-by-cohort interactions with a single coverage $\times$ cohort linear trend, the form any pre-existing differential trend would take, yields a slope indistinguishable from zero ($+0.00001$ per cohort year, $t = 0.01$, for men; -0.0011, $t = -1.34$, for women, a test made conservative by including the lightly exposed 1989–1990 cohorts). Even taking the women's insignificant point estimate at face value and extrapolating it forward to the 1994 cohort would deliver less than a tenth of the estimated break: there is no latent pre-rollout trend for the exposed cohorts to be continuing. Second, the result is not an artifact of *where* men are observed rather than where they were formed. Re-estimating the design at the state level, where the treatment is the population-weighted formative 4G exposure of a cohort's state and movement across commuting zones within a state cannot bias it, reproduces the break: the state-level currently-married estimate is -0.035 ($t = -2.1$), negative and significant as at the commuting-zone level. Because the large majority of young-adult moves are within state, this addresses the bulk of the concern that men sort across places in ways correlated with marriage; the residual it cannot reach is cross-state migration. Assigning exposure by state of birth, which no adult move can affect, and re-estimating on the US-born (84 percent of the sample) with birth-state and cohort fixed effects, the estimates remain negative throughout: the women's break survives at conventional significance (-0.037 currently married, -0.043 ever married), while the men's point estimate attenuates to -0.014 ($t = -1.22$), as a birth-state assignment must when it replaces the formative-years commuting zone for the many whose families moved. The men's birthplace confidence interval contains the state-level estimate of -0.035, so the strictest design is consistent with, though underpowered to confirm, the migration-robust result; what it rules out is that cross-state migration produced the break. Third, inference is conservative: with roughly 490 commuting-zone clusters conventional clustering is already reliable, and the within-cohort high-versus-low-exposure contrast for men additionally survives a wild-cluster bootstrap on the exposed-and-adjacent cohorts (474 clusters; $p = 0.020$).

Fourth, and most pointedly, the break is not the Great Recession. The exposed cohorts entered the labor market during and after the 2008–2010 downturn, raising the concern that the cohort pattern reflects local labor markets rather than mobile internet. For each cohort and zone we construct the mean local prime-age (25–54) unemployment rate over the cohort's labor-market-entry years, ages 18–25 (an analogous cohort-by-zone construction, from the ACS) and add it to equation (24). The exposed-cohort coefficients are essentially unchanged: the 1994 estimate moves from -0.040 to -0.038 for men and is unmoved for women, and the joint test remains $p \leq 0.0001$ for both. The control itself enters with a *positive* sign for men ($+0.26$, $t = 2.1$): cohorts forming in weaker local economies married somewhat more, not less, so the recession works against the estimate rather than generating it. The break is tied to the timing of the rollout, not of the downturn.

Fifth, the break does not hinge on the ever-married definition: Figure 4 re-estimates the men's path on the currently-married rate, which also reflects divorce rather than only first marriage, and finds the same flat-then-negative shape, somewhat larger in magnitude (reaching -0.057 at the most-exposed cohort), and with its pre-rollout coefficients likewise jointly flat ($F = 0.96$, $p = 0.44$).

Sixth, the treatment is a continuous dose, and the modern difference-in-differences literature cautions that in such designs the late coefficients are identified off intensity contrasts once no unexposed units remain (de Chaisemartin and D'Haultfœuille, 2020; Callaway, Goodman-Bacon, and Sant'Anna, 2024). That caution has less purchase here than the design's continuous treatment might suggest, because untreated units never disappear from the comparison: over a quarter of the estimation sample (a majority of commuting zones unweighted) sits in zones with literally zero 2013 4G coverage, so for every cohort, including the youngest, the β_k are anchored by a genuine exposed-versus-unexposed contrast rather than by intensity differences alone. For cohorts born through 1989 the contrast is degenerate in the other direction, every zone carrying exactly zero formative exposure; from 1990 onward the coefficients blend the zero-coverage anchor with dose variation across covered zones. The zone fixed effects difference out level selection into coverage, but not, by themselves, heterogeneous responses across the dose distribution. This is why Table 8 complements the continuous path with a transparent binary contrast: splitting zones at the median 2013 coverage and comparing exposed cohorts across the two groups reproduces the negative marriage response ($p = 0.020$, wild-cluster bootstrap), an estimate whose identifying variation is exactly the high-versus-low comparison the concern calls for. Displays of the coverage distribution across zones and its evolution within zones over the rollout accompany the replication materials.

Finally, the near-parallel female path is itself a check the design was not built to pass: a

TABLE 8. Robustness of the cohort marriage estimate

Specification	Men	Women
Baseline event study (1994 cohort)	-0.040^{***}	-0.066^{***}
+ local unemployment control	-0.038^{***}	-0.066^{***}
reference cohort 1985	-0.043^{***}	-0.066^{***}
State-level treatment, currently married	-0.035^{**}	-0.053^{***}
Birthplace treatment, currently married	-0.014	-0.037^{**}
Birthplace treatment, ever married	-0.006	-0.043^{***}
Pre-rollout placebo: coverage $\times$ cohort trend	$+0.000$	-0.001
Wild-cluster bootstrap, high vs. low	-0.015^{**}	–

Notes. Variants of the cohort design of Section VI, on all men and women aged 25–34 (all education levels); entries are the marriage coefficient. The event-study rows report the 1994 (most-exposed clean) cohort coefficient on the ever-married rate from equation (24) (cluster-robust $t = -4.5$ men, -7.6 women; pre-trend joint $p = 0.78/0.79$), the second row adding formation-window local unemployment, which itself enters positively ($+0.26$, $t = 2.1$), and the third repeating the baseline with the 1985 pre-rollout cohort as the omitted reference (pre-rollout joint $p = 0.65$ men, 0.69 women). The pre-rollout placebo row restricts the sample to the 1983–1990 cohorts and fits a single coverage $\times$ cohort linear trend in place of the cohort interactions, the form a break predating the rollout would take; entries are the trend slope per cohort year ($t = 0.01$ men, -1.34 women), and including the lightly exposed 1989–1990 cohorts makes the test conservative. The state-level row aggregates treatment to the state and uses the currently-married rate, so within-state migration cannot bias it ($t = -2.1$ for men, -2.7 for women). The birthplace rows reassign treatment by state of birth (US-born, 84 percent of the sample; birth-state and cohort fixed effects), which no adult move can affect. The bootstrap row is a wild-cluster bootstrap-t (Rademacher weights, 474 commuting-zone clusters) on the within-cohort high-versus-low-exposure binary ($p = 0.020$); this within-cohort binary is a men-specific check; the women's break is identified off the cross-cohort path (Figure 3), where the within-cohort contrast is near zero, so the row does not apply for women. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: authors' calculations from ACS via IPUMS and FCC mobile-deployment data.

generic place–cohort shock to marriage need not respect the two-sidedness of the marriage market, whereas a genuine decline in partnership formation must appear, as it does, on both sides at once.

VI.F The break by education

The paper's labor story concerns non-college educated men, so we re-estimate the cohort design directly on that group and, for contrast, on college men and on women of each schooling, splitting the cells by education. Figures 5–8 report the paths.

For non-college educated men (Figure 5) the break sits on the margin the labor story uses. The currently-married rate, the "married share" of Section III, is flat across the pre-rollout cohorts and falls to -0.036 at the most-exposed clean cohort (1994; the 1995–

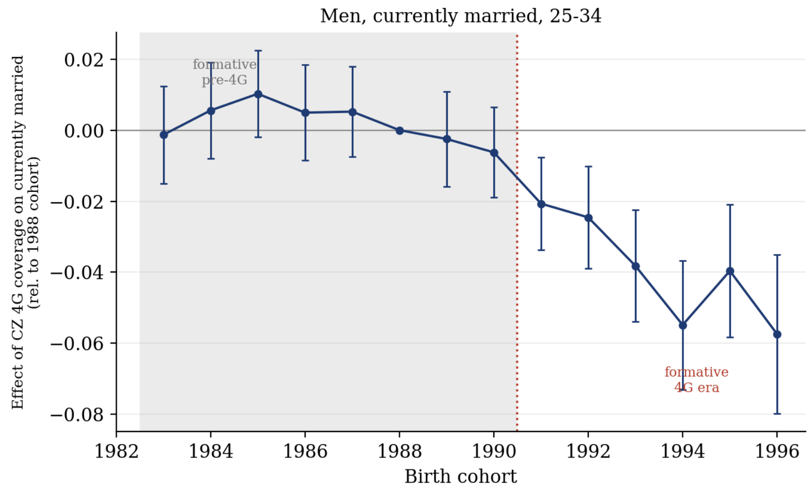

FIGURE 4. Cohort event study for men (all education levels, ages 25–34), currently-married rate, the non-redundant panel of the marriage-definition check. The break is the same flat-then-negative shape as the ever-married path of Figure 3, somewhat larger in magnitude. Shaded pre-rollout region and 95 percent confidence bands as in Figure 3.

96 cohorts are observed over a compressed age range), with the exposed cohorts jointly significant ($p = 0.006$); the ever-married rate moves in the same direction but does not itself reject at conventional levels in this split ($p = 0.27$). We flag the tension rather than gloss it: the mechanism is about first marriages that fail to form, the margin the ever-married rate captures, while the sharper break here sits on the currently-married rate, which also reflects divorce. We read the gap as power rather than economics: the two point estimates move together on the same cohorts, the pooled ever-married path of Figure 3 rejects decisively, and splitting by education thins the cells; we do not rest the case on the currently-married coefficient alone. Concretely, the split's ever-married coefficient at the most-exposed cohort carries a standard error near 0.010, giving a minimum detectable effect of about 0.028 at conventional power (eighty percent, five percent two-sided), larger than its own 0.020 point estimate and comparable to the 0.036 currently-married break, so the non-rejection reflects the thinned split's limited power rather than an absent effect. For college men (Figure 6) neither margin breaks (the 1994 point is -0.021, wide-banded; jointly $p = 0.66$): the effect sits in exactly the group whose participation fell. The split also explains why the all-sample estimate of Figure 3 runs a touch larger than the within-

Non-college men: marriage by cohort (early-4G vs late-4G zones)

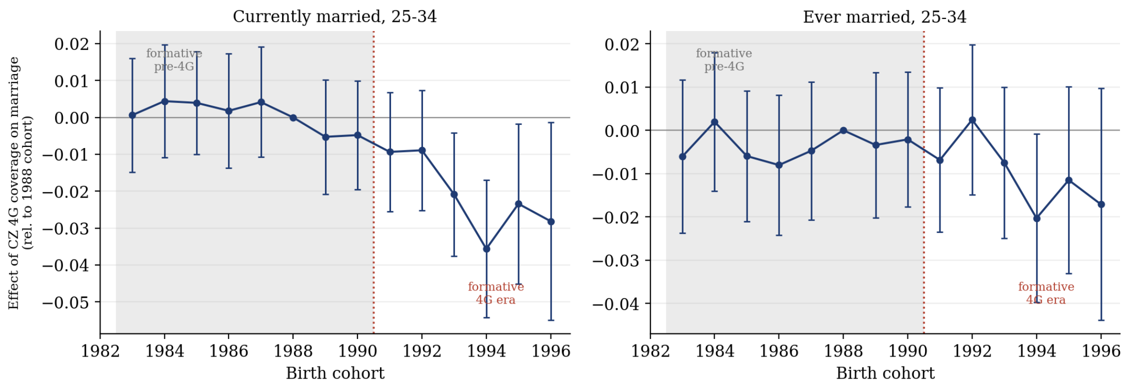

FIGURE 5. Cohort event study, non-college educated men. Left: currently married; right: ever married. Shaded pre-rollout region and 95 percent confidence bands as in Figure 3.

non-college effect: high-coverage cohorts skew slightly more educated, and college men marry later, so the pooled coefficient carries a sliver of education composition that the split removes.

For women the break appears for *both* schooling groups (Figures 7 and 8): at the same 1994 cohort non-college educated women fall -0.050 on the currently-married rate and college women -0.064, with the exposed cohorts jointly significant ($p < 0.001$ and $p < 0.05$); the averages across all exposed cohorts, -0.024 and -0.043, tell the same story. These education splits are reported on the currently-married rate, matching the men's split above; on the ever-married rate, the margin of the pooled headline in Figure 3, the women's break likewise appears at both schooling levels. On either margin the marriage response is thus both broader among women, appearing at both schooling levels, and at least as deep as the non-college-specific male pattern. This is the cross-gender symmetry of Figure 3 seen up close, and it closes the loop with the accounting: the non-college educated women whose participation *rose* (Section III) are the same women whose marriage fell.

VI.G What this design can and cannot show

We are explicit about the limits of this evidence. The treatment is local 4G coverage, which is not randomly assigned; the within-cohort design absorbs every fixed feature of

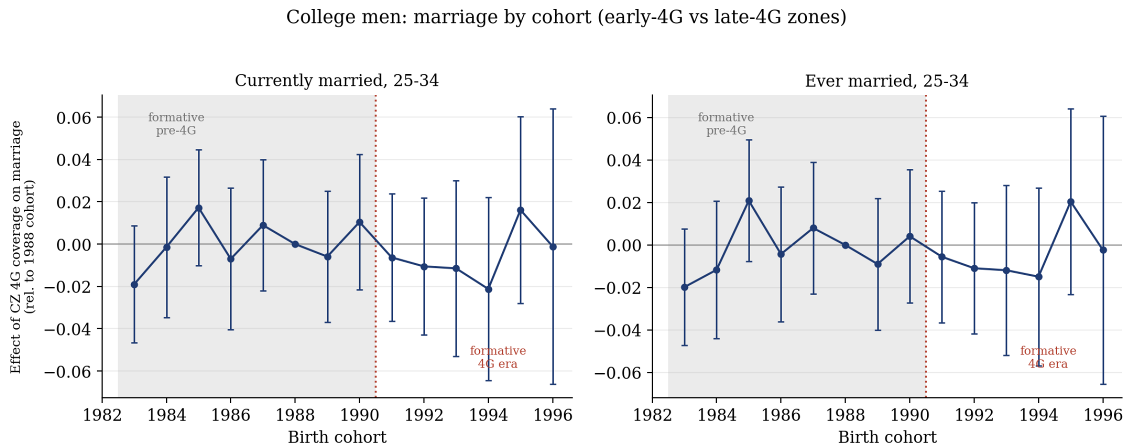

FIGURE 6. Cohort event study, college men. Left: currently married; right: ever married. Neither margin breaks. Shaded pre-rollout region and 95 percent confidence bands as in Figure 3.

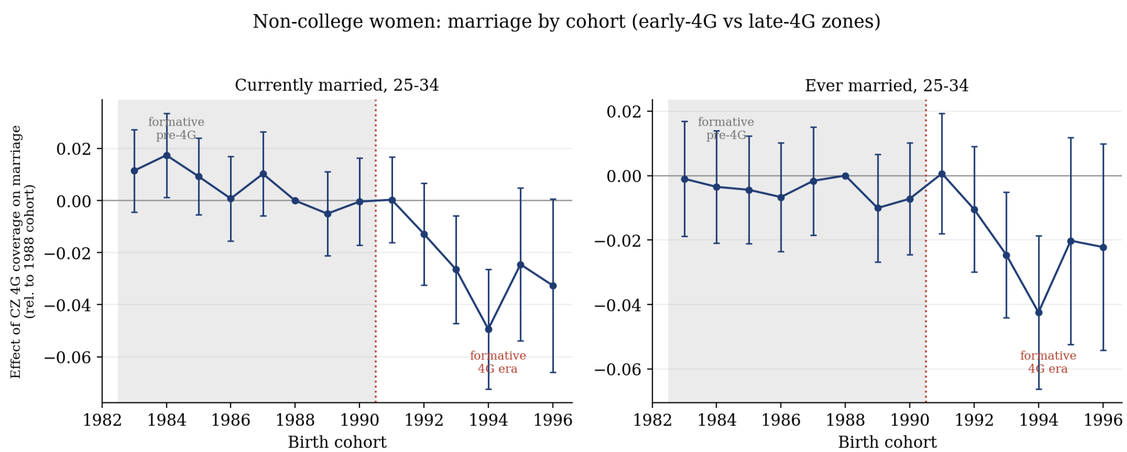

FIGURE 7. Cohort event study, non-college educated women. Left: currently married; right: ever married. Shaded pre-rollout region and 95 percent confidence bands as in Figure 3.

a place, and the flat pre-rollout path rules out a smooth place-specific marriage trend of the kind that would otherwise confound the comparison, but no design of this kind can exclude every cohort-specific shock that happened to track the rollout's timing within zones. This is why we read the cohort evidence as consistent with a digital origin for the marriage decline rather than as a stand-alone causal magnitude. The composition accounting of Section III and the cross-gender redistribution require no instrument and no parallel-trends assumption. The within/between decomposition is an accounting identity; reading its between component as the labor-market consequence of the marriage decline,

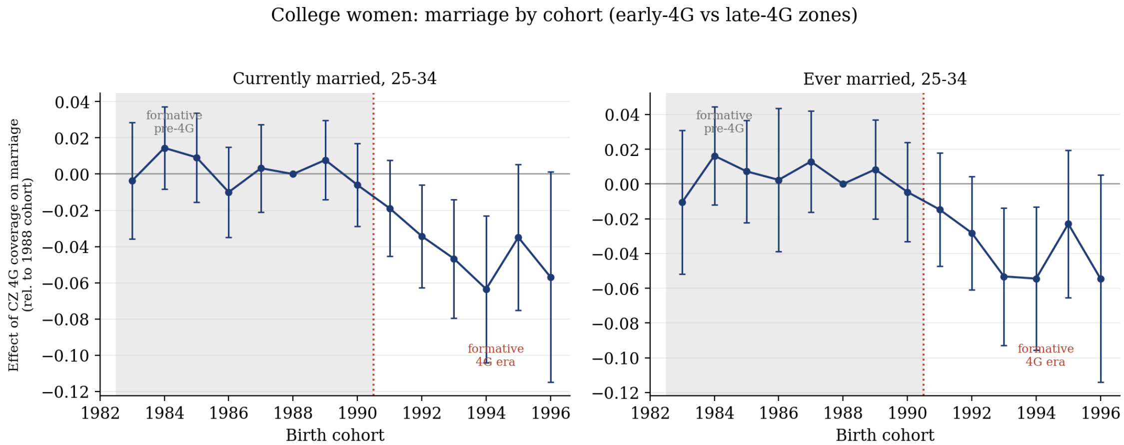

FIGURE 8. Cohort event study, college women. Left: currently married; right: ever married. Shaded pre-rollout region and 95 percent confidence bands as in Figure 3.

rather than a relabeling of fixed types, rests on the anti-selection evidence of Section III: the cross-gender symmetry of the aggregates and the rise in total work for both married groups, neither of which a pure relabeling would produce. Given that evidence, the married–single participation gap delivers the marriage decline's labor-market consequence through composition, and the opposite-signed movement of the two genders' aggregates is itself a fact about the data. The cohort event study tells us that the marriage decline these accounting facts operate on is, in its timing and its cross-gender symmetry, what a digital cause would produce.

VI.H From marriage to labor: the size of the compositional channel

Whatever its cause, the marriage shift has a labor consequence through composition, and its size follows from the decomposition. Married non-college educated men participate 8.3 points above single ones at the start of the window and 10.9 points above by its end. The 10.6-point decline in their married share, multiplied by this gap, is exactly the between component of Section III: 0.9 points at the base-period gap, 1.2 points at the end-period gap. The marriage shift is thus large enough to account for the entire composition component, about a third of the aggregate participation decline, without appeal to any additional channel. Because participation is measured at the same young ages where the marriage shift appears, a *delay* in marriage, not only a permanent decline in who ever marries, is sufficient to move this channel. We flag rather than lean on the fact that marriage has been declining for decades, so attributing the entire 2005–19 compositional shift to the recent rollout would overstate the case; the conservative reading is that the digitally driven portion of the marriage decline accounts for a corresponding portion of the composition component, itself about a third of the aggregate participation decline.

VII. CONCLUSION

The post-2007 marginal change in young non-college educated men's labor supply was social in origin rather than occupational. The composition of these men shifted, about ten points from married to single between the mid-2000s and late-2010s windows, a long-run marriage decline whose digital timing shows up in cohorts rather than in the calendar, and this shift, given the participation gap between the cells, accounts for roughly a third of the aggregate decline, while the within-group rate of the single continued its longer-run slide without a smartphone-era break. The same marriage-market event, decomposed identically for young women, enters with the opposite sign and accounts for the majority of their 3.2-point aggregate rise: in rate terms the digital-era marriage decline redistributed young non-college labor supply across the genders rather than reducing it. Time-diary data show this redistribution is one of measurement, not of effort: married men and women both supply more total work, market plus home production, than their single counterparts, but only men's increment falls on the counted market margin. The behavioral residual left after composition is specifically male, and the time-diary record shows where the technology's first-order effects landed: half an hour per day of additional solitary screen leisure, against five minutes for the comparison group, and twenty minutes per day less with friends. That screen time came out of other leisure and the company of others, not out of work: in the cell where screens rose most, market work held flat, so the decline in young men's work hours that has been read as a substitution of screens for work is, on this evidence, substantially the same compositional shift that moves participation.

The model places this account on a single primitive: endogenizing labor supply in the household-production framework of Hudson and Moscoso Boedo (2026b), the gross complementarity between the digital bundle and time that lowers relational-capital investment also makes the household a deeper relationship supports raise the marginal value of market income, so that relational capital and market work move together. Solved as a transition along the same falling phone price that paper uses to explain fertility, with nothing fit to labor, the model reproduces the observed fertility decline as an untargeted prediction and, through composition, traces a participation decline of the size the decomposition measures. A within-place comparison of birth cohorts shows the cohorts who came of age on 4G married measurably less, with a flat pre-rollout path that breaks at the exposed cohorts, tracks the timing of the rollout rather than of the recession, appears in parallel for women, and is concentrated among non-college educated men; the paper's weight rests on the accounting. The decline did not operate by pulling men who stayed single out of work; it operated through the marriages they did not form, and the jobs followed the marriages. The redistribution of that shift across the genders is, we think, the fact about the digital era's labor consequences that this literature has most overlooked.

References

Abraham, Katharine G., and Melissa S. Kearney. 2020. "Explaining the Decline in the US Employment-to-Population Ratio: A Review of the Evidence." *Journal of Economic Literature* 58 (3): 585–643.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst. 2021. "Leisure Luxuries and the Labor Supply of Young Men." *Journal of Political Economy* 129 (2): 337–382.

Ahituv, Avner, and Robert I. Lerman. 2007. "How Do Marital Status, Work Effort, and Wage Rates Interact?" *Demography* 44 (3): 623–647.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015. "The Skill Complementarity of Broadband Internet." *Quarterly Journal of Economics* 130 (4): 1781–1824.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020. “The Welfare Effects of Social Media.” *American Economic Review* 110 (3): 629–676.

Atasoy, Hilal. 2013. “The Effects of Broadband Internet Expansion on Labor Market Outcomes.” *ILR Review* 66 (2): 315–345.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103 (6): 2121–2168.

Bhuller, Manudeep, Andreas R. Kostøl, and Trond C. Vigtel. 2020. “How Broadband Internet Affects Labor Market Matching.” IZA Discussion Paper 12895.

Binder, Ariel J., and John Bound. 2019. “The Declining Labor Market Prospects of Less-Educated Men.” *Journal of Economic Perspectives* 33 (2): 163–190.

Braghieri, Luca, Ro’ee Levy, and Alexey Makarin. 2022. “Social Media and Mental Health.” *American Economic Review* 112 (11): 3660–3693.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant’Anna. 2024. “Difference-in-Differences with a Continuous Treatment.” NBER Working Paper 32117.

de Chaisemartin, Clément, and Xavier D’Haultfœuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *American Economic Review* 110 (9): 2964–2996.

Dettling, Lisa J. 2017. “Broadband in the Labor Market: The Impact of Residential High-Speed Internet on Married Women’s Labor Force Participation.” *ILR Review* 70 (2): 451–482.

Eberstadt, Nicholas. 2016. *Men Without Work: America’s Invisible Crisis*. West Conshohocken, PA: Templeton Press.

Federal Communications Commission. 2020. *Mobile Deployment Form 477 Data*, 2010–2019. Washington, DC: Federal Communications Commission.

Flood, Sarah M., Liana C. Sayer, and Daniel Backman. 2023. *American Time Use Survey Data Extract Builder: Version 3.2* [dataset]. College Park, MD: University of Maryland and Minneapolis, MN: IPUMS. `https://doi.org/10.18128/D060.V3.2`.

Hjort, Jonas, and Jonas Poulsen. 2019. “The Arrival of Fast Internet and Employment in Africa.” *American Economic Review* 109 (3): 1032–1079.

Hudson, Nathan, and Hernán Moscoso Boedo. 2026a. "The Collapse of Teen Fertility in the Digital Era." Working paper, University of Cincinnati.

Hudson, Nathan, and Hernán Moscoso Boedo. 2026b. "Wide and Shallow: Digital Technology and the Post-2007 Fertility Decline." Working paper, University of Cincinnati.

Kimbrough, Gray. 2020. "Xboxes and Ex-workers? Gaming and Labor Supply of Young Adults in the U.S." Working paper.

Korenman, Sanders, and David Neumark. 1991. "Does Marriage Really Make Men More Productive?" *Journal of Human Resources* 26 (2): 282–307.

Krueger, Alan B. 2017. "Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate." *Brookings Papers on Economic Activity* (Fall): 1–87.

Myers, Caitlin K., and Ezekiel Hooper. 2026. "Is the iPhone Birth Control? Causal Evidence from AT&T's 2007–2011 Carrier Monopoly." NBER Working Paper 35310.

Rosenfeld, Michael J., Reuben J. Thomas, and Sonia Hausen. 2019. "Disintermediating Your Friends: How Online Dating in the United States Displaces Other Ways of Meeting." *Proceedings of the National Academy of Sciences* 116 (36): 17753–17758.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rodgers, Jonathan Schroeder, and Kari C. W. Williams. 2024. *IPUMS USA: Version 15.0* [dataset]. Minneapolis, MN: IPUMS. `https://doi.org/10.18128/D010.V15.0`.

Tolbert, Charles M., and Molly Sizer. 1996. "U.S. Commuting Zones and Labor Market Areas: A 1990 Update." Staff Paper 9614. Washington, DC: USDA Economic Research Service.

Zuo, George W. 2021. "Wired and Hired: Employment Effects of Subsidized Broadband Internet for Low-Income Americans." *American Economic Journal: Economic Policy* 13 (3): 447–482.