AI adoption, local labour demand, and marriage

Claudio Costanzo*

September 16, 2026

Abstract

Earlier waves of automation reduced marriage in the local labour markets they hit. This paper shows that the diffusion of artificial intelligence did the opposite over 2010–2021, and identifies the reason. I measure local exposure with a shift-share that weights industry AI adoption from the Annual Business Survey Technology Module by baseline commuting-zone employment shares, instrumented with its European analogue. A one-standard-deviation rise in exposure raises the share of adults in a first marriage by 1.3 percentage points and lowers the divorced-to-married ratio by 1.9 points. A matching model of the local marriage market shows that a sectoral shock can raise the marital surplus through the level of earnings or through their variance, that the two channels leave different signatures in the data, and that a shock of given size raises marriage by more the more of it accrues to men. The data select the earnings channel: exposure raises employment for both sexes, draws employment into the most AI-exposed occupations, widens the male–female earnings gap in pay per hour and not only in hours worked, and adds marriages in proportion to the couple's exposure, with none of the tilt toward mixed couples that risk pooling requires. AI adoption arrived as a local labour-demand expansion whose incidence favoured men, the mirror image of the robot and trade shocks that lowered men's relative position. The same exposure predicts nothing over 2000–2007, and the effects build through the 2010s, visible well before the pandemic.

Keywords: Artificial intelligence; labour demand; marriage; divorce; assortative matching; shift-share.

JEL Codes: J12; J31; J24; O33; R23

*ECARES, Université Libre De Bruxelles. Contact: claudio.costanzo@ulb.be

1 Introduction

Artificial intelligence has emerged as the defining general-purpose technology of the past decade. Driven by increases in computing power, the availability of large training datasets, and breakthroughs in deep neural networks, AI is reshaping the economic frontier (Agrawal et al., 2019). Unlike earlier waves of computerisation and robotics, which automated routine manual and clerical tasks, modern AI applications are directed at non-routine cognitive and interactive functions (Webb, 2019; Felten et al., 2021; Acemoglu et al., 2022). Technological shocks of this magnitude rarely remain confined to the workplace, and one of the margins along which local labour markets adjust is the household.

The existing evidence on technology and family formation points in one direction. Exposure to industrial robots reduces marriage and raises divorce in affected commuting zones (Anelli et al., 2024), and the decline of manufacturing employment lowers the marriage-market value of young men and, with it, the marriage rate (Autor et al., 2019). The mechanism in that literature is well understood: these shocks fell on male-intensive industries, they lowered men's relative earnings and employment, and fewer marriages followed (Wilson, 1996).

This paper asks what the AI wave did, and finds the opposite sign. The technologies in question are the machine learning, machine vision, natural language, voice and automated guided vehicle tools deployed by 2020, the wave that preceded generative AI. Commuting zones more specialised in the industries that went on to adopt them experienced relatively more first marriages and relatively fewer divorces between 2010 and 2021. The reason is that AI adoption, as measured by firms' reported use of these technologies, accompanied an expansion of local labour demand in the places it reached, and the incidence of that expansion favoured men.

Two descriptive facts anchor that reading. First, corporate AI activity was negligible before 2010 and surged immediately after. Second, over the decade of diffusion, while marriage continued its long national decline, the share of adults in couples where both partners work in highly AI-exposed occupations turned upward from about 2013 and ended the period above its starting level, the only couple type to do so. Marriage thus rose among the workers most exposed to the new technology in the very years it spread, the reverse of the historical pattern.

I then build a model of a local labour market and of the marriage market it supports. A national pattern of industry adoption moves the local wage, a price that clears the labour market of the commuting zone as a whole, so the shock reaches households whose own industries adopted nothing. How much of it goes to men and how much to women is left to be measured, since the division depends on which industries carry the technology. The division bears on marriage because the household converts the higher earner's wage into consumption at the higher rate, once home time is allocated efficiently and men earn more at baseline, so a shock of given size raises the marital surplus by more the more of it men capture. Marriages, divorces and singles are then equilibrium quantities in a matching market with transfers and an outside option, so the sex ratio and the participation margin both shape the equilibrium response. Within that framework a sectoral shock can raise the surplus either through the level of earnings or through their variance, and the two channels carry different and testable signatures. The risk channel requires exposure to raise earnings dispersion, to displace workers from exposed occupations, and to place the marriages it creates in couples whose partners are differently exposed. The labour-demand channel requires exposure to raise employment and men's relative earnings, and places its marriages in proportion to how exposed both partners are. Dispersion is the weakest of the three tests, because a demand expansion that raises pay at the top of a local distribution widens it as readily as an increase in risk does. The other two requirements discriminate between the channels.

The empirical strategy relates 2010–2021 long differences in marital and labour-market outcomes to a Bartik-style measure of local AI exposure, which I construct by weighting national industry adoption rates from the Annual Business Survey (ABS) Technology Module by baseline 2010 commuting-zone industry shares and instrument with the analogous measure built from European industry adoption.

In the fully controlled two-stage least squares specifications, a one-standard-deviation increase in exposure raises the share of adults in a first marriage by 1.30 percentage points and lowers the divorced-to-married ratio by 1.87 points, both significant at the 1 percent level, against sample-mean changes of -1.18 and -0.49 points. The same increase raises the employment rate by 2.09 percentage points and widens the male–female earnings gap, while raising the share of employment in the occupations the Webb (2019) index identifies as most exposed to replacement. Over this decade, that is, the local labour markets that adopted AI most intensively employed relatively more of their AI-exposed workers. The dispersion of annual earnings is statistically unchanged. Measured as a wage per hour it widens, but a demand expansion produces that widening as readily as rising risk does, so the movement does not discriminate between the channels.

Two features of the timing support a causal reading. First, applying the identical exposure measure to 2000–2007, before AI diffusion had begun, produces nothing on any outcome. Second, tracing the coefficient by endpoint year shows the labour-demand and earnings-gap effects building steadily from 2012 onward instead of appearing abruptly with the pandemic. One confounder lies between the placebo window and the estimation window. The same commuting zones were hit unusually hard by the Great Recession, so part of the post-2010 expansion is recovery, and conditioning on each commuting zone's own 2007–2010 experience attenuates the estimates without eliminating them.

The identifying variation comes from 46 industry-level shocks whose Borusyak et al. (2023) effective number is 13.5. When I reshuffle the adoption vector at random across industries and re-estimate, the true industry ranking is not extreme, because the commuting-zone share matrix has few effective dimensions and almost any assignment of these numbers generates a similar geographic map. The estimand is therefore the effect of specialising in the industries that went on to adopt AI, not the effect of AI holding constant everything correlated with adoption. The diagnostics are reported in full in an appendix. A comparison of industries within the same commuting zone shows, moreover, that workers whose own industries adopted AI exhibit no differential response, so the effect is a local general-equilibrium one, exactly as a labour-demand interpretation implies.

This paper contributes to the literature on how technology-driven shifts in labour demand shape household behaviour. Alongside the marriage-market evidence above, that literature traces such shocks into fertility and its timing over the life cycle (Costanzo, 2025; Matysiak et al., 2023; Schaller, 2016), into the labour supply of women and the division of paid and unpaid work within couples (Greenwood and Guner, 2004; Greenwood et al., 2005; Ngai and Petrongolo, 2017; Black and Spitz-Oener, 2010; Aguiar and Hurst, 2007), and into the household's role in insuring the earnings risk such shocks create (Kotlikoff and Spivak, 1981; Rosenzweig and Stark, 1989; Hess, 2004). In each case the shock reaches the household through the labour market, so its effect depends on whose wages and employment it moves; the present paper exploits the same route.

Against that background, the contribution is to provide the first evidence linking AI diffusion to marital outcomes, and to show that the sign of the demographic response is not a property of the technology but of the labour-demand shock that accompanies it and of that shock's incidence by sex.

Section 2 documents the stylized facts, Section 3 develops the model, Section 4 describes the data and the empirical strategy, Section 5 presents the results, and Section 6 concludes.

2 Stylized facts

This section documents two descriptive facts that motivate the design. Occupational exposure to AI is measured throughout with the Webb (2019) index. Webb identifies the patents that pertain to artificial intelligence, reduces each patent title to its verb–object pairs, and does the same to the O*NET task statements that describe what each occupation does. A task is exposed to the extent that its verb–object pair recurs in AI patents; an occupation's score aggregates that overlap across its tasks; and the score is expressed as a percentile of the employment-weighted distribution across occupations, which is the form in which the index enters everything below.

1. AI innovation has surged after 2010. To benchmark the onset of the modern AI boom, I track the release of landmark artificial intelligence systems over time. While academic research into neural networks spans decades, the transition of the technology into a commercial force is recent. Figure 1 plots the annual count of "notable" AI machine learning models developed by the corporate sector between 2000 and 2024, using data curated by Epoch AI. Prior to 2010 the release of notable AI systems by industry was near zero, and development was confined almost entirely to academic research. Immediately thereafter the corporate sector experienced an exponential takeoff in model deployment.

This 2010 pivot aligns with the breakthrough in deep convolutional neural networks that demonstrated the first significant technical advantage of deep learning over rule-based algorithms (Krizhevsky et al., 2012), and with the subsequent surge in firm-level investment and skill demand (Babina et al., 2024). Appendix Figure C1 shows the same acceleration in AI patents granted by the USPTO. The timing motivates 2010 as the baseline year, so that local industry shares are fixed before diffusion could plausibly have induced labour reallocation.

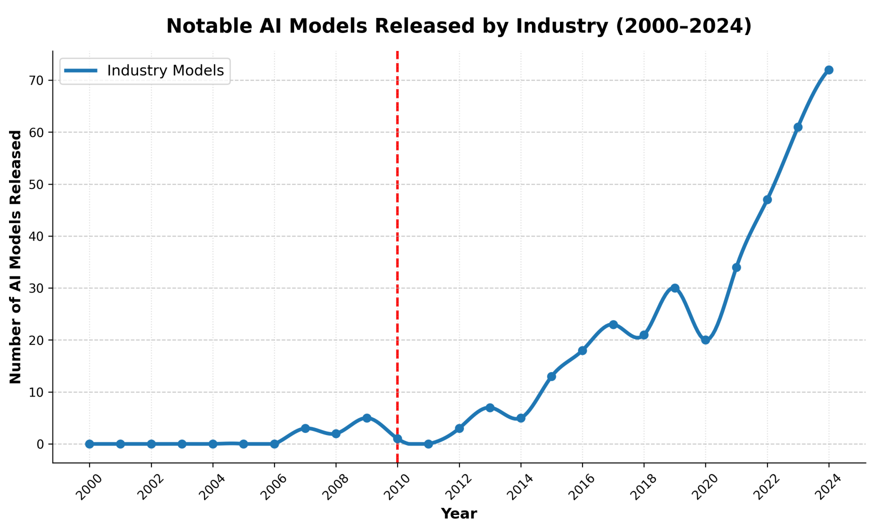

Figure 1: **Notable AI Models Released by Industry (2000–2024).** The smoothed line represents the annual count of landmark AI machine learning models developed by the corporate sector. The vertical dashed line marks the 2010 structural break, illustrating the transition of AI from academic research to corporate investment. Source: Epoch AI, *Notable AI Models Dataset*.

2. Marriage patterns diverge along a gradient of AI exposure after 2010. The second fact uses ACS microdata (Ruggles et al., 2025) to track the evolution of marriage by the joint AI exposure of the couple. Starting from person records, I link each respondent to a spouse within the household. Couples are classified into four mutually exclusive types using the Webb (2019) percentile cutoffs: low–low (both $\leq$ 33rd), high–high (both $\geq$ 66th), extreme mixed (one $\leq$ 33rd, one $\geq$ 66th), and any–medium (at least one spouse strictly between the 33rd and 66th percentiles, serving as a broad population benchmark). Throughout the paper a "couple" means a married couple observed with both spouses present in the household, so couples and marriages are used interchangeably.

To account for endogenous labour force participation, whereby technology shocks might shift workers into non-employment, the denominator is the entire working-age population aged 18–64, regardless of employment status. The universe is the household population throughout. The ACS added group quarters in 2006, and because only about nine percent of that population is married, including it would produce a mechanical break in every couple-type share between 2005 and 2006. No estimation window straddles that break (group quarters are a stable 3.1 to 3.2 percent of the 18–64 population in 2000, 2007, 2010, 2019 and 2021 alike), so the restriction bears only on this figure.

Figure 2 plots the three focal couple types in levels in the left panel and as changes since 2002, in percentage points, in the middle panel. Both are shown because a share normalised to its own base year can suggest a much larger movement than the data contain; in levels, the low–low decline is a fall of roughly one and a half percentage points of the population. The any–medium benchmark appears separately in the right panel. It is a residual category (at least one spouse anywhere in the middle third of the exposure distribution), so it covers roughly a quarter of the population against three to nine percent for each of the other three, and on a shared axis it compresses them into an unreadable band.

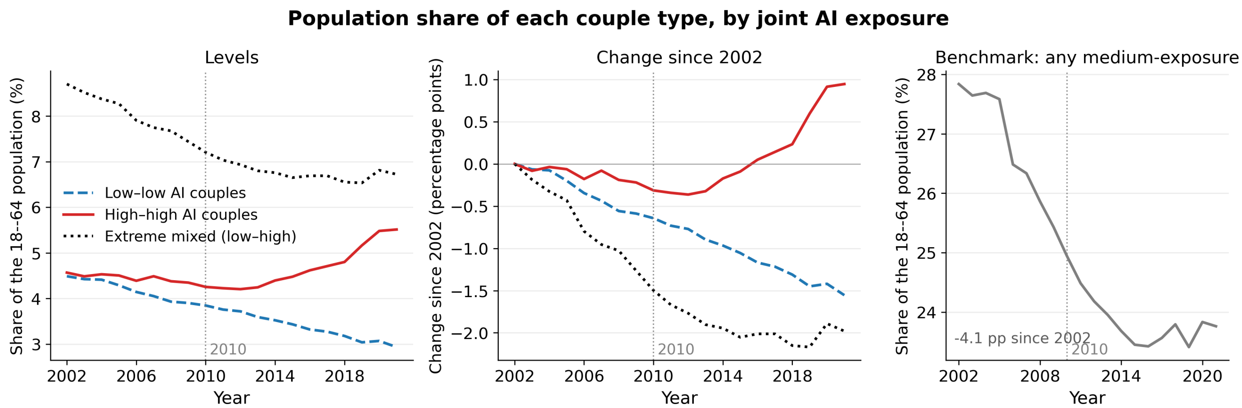

Figure 2: **Population share of each couple type, by joint AI exposure.** Left panel: share of the 18–64 household population in each of the three focal couple types. Middle panel: change in that share since 2002, in percentage points. Right panel: the any–medium benchmark on its own scale, since it covers roughly a quarter of the population and would otherwise compress the other three into an unreadable band. Couple types use the Webb (2019) percentile cutoffs described in the text. The denominator includes all individuals aged 18–64, employed or not, so the series are not driven by movements in labour force participation, and is restricted to the household population so that the 2006 introduction of group quarters into the ACS does not induce a break. Occupations are observed contemporaneously in this figure; the regressions in Section 5.3 instead freeze both partners' occupations at 2010. Source: Ruggles et al. (2025) and Webb (2019).

The data reveal a gradient corresponding to the couple's joint exposure to AI. The population share of low–low couples declines throughout the sample period, consistent with the broader national decline in marriage, and extreme mixed couples follow the same secular path, falling further and faster. High–high couples decline too through the first decade, though at about half the low–low pace, falling 0.36 percentage points to a trough in 2012 against 0.77. From 2013 the share turns upward, and by the end of the period it stands about one percentage point above its 2002 level, the only one of the four types to recover its starting share. The gradient is modest in absolute terms but consistent in ordering. It serves as motivation only, since occupations here are observed contemporaneously and composition therefore shifts as workers move between occupations.

Because AI exposure is positively correlated with educational attainment, the gradient could reflect a rising college premium in the marriage market, with occupational exposure merely proxying for education. Appendix Figure C2 replicates the exercise restricting both numerator and denominator to individuals with a Bachelor's degree or higher. The ordering of the series is unchanged, which suggests that educational attainment alone does not generate the gradient.

3 Theoretical Framework

This section builds a model of a local labour market and of the marriage market that market supports. The model has three tasks. First, it must let a sectoral technology shock reach households whose own industries adopted nothing: within a commuting zone workers move across industries, so adopting and non-adopting employers compete for the same labour, and a demand shift anywhere arrives at every household as a movement in the local wage. Second, it must let the incidence of the shock differ by sex, because a shock that raises men's earnings and a shock of the same size that raises women's are not the same event in a marriage market. Third, it must keep both candidate channels open, a movement in the level of earnings and a movement in their variance, so that the choice between them is left to the data. The three blocks below accomplish these tasks in turn: local labour demand fixes wages and employment, a household technology turns a pair of wages into a marital surplus, and a matching market with an outside option turns surpluses into marriages, divorces and singles.

Local labour demand. Commuting zone c contains J industries with baseline employment shares s_{cj}. Industry j adopts AI at the national rate a_j, which shifts its labour demand,

$$\ln L^d_{cj} = d_{cj} + a_j - \eta \ln w_c, \tag{1}$$

with $\eta > 0$ the elasticity of labour demand and w_c the local wage. Labour is mobile across industries within the commuting zone, so w_c is common to all of them, and imperfectly mobile across commuting zones, so local supply $\ln L^s_c = \ell_c + \varepsilon \ln w_c$ slopes upward with $\varepsilon < \infty$. Clearing the local market and aggregating with baseline shares,

$$\Delta \ln w_c = \frac{1}{\eta + \varepsilon} \sum_j s_{cj} a_j \;\equiv\; \phi\, B_c, \qquad \Delta \ln L_c = \varepsilon\, \phi\, B_c, \tag{2}$$

where $B_c \equiv \sum_j s_{cj} a_j$ is the shift–share exposure constructed in Section 4.1. Within the model the shift–share is the local wage shifter itself, since initial industrial composition determines how far local labour demand moves (Bartik, 1991; Moretti, 2011).

Equation (2) already settles the level at which the shock operates. The wage is a price that clears the labour market of the commuting zone as a whole. Two workers in the same zone, one in an industry that adopted AI intensively and one in an industry that adopted nothing, face the same w_c, and nothing in the model makes the first respond differently from the second. That is a prediction Section 5.3 tests against the alternative that the effect runs through the adopting industries themselves.

The incidence of the shock across the sexes remains to be determined. Men and women are imperfect substitutes in production, and industries differ in how much of local male and of local female employment they account for. Writing m_{cj} and f_{cj} for those two shares, the shock moves relative demand for male labour by $\sum_j (m_{cj} - f_{cj})a_j$, so that

$$\Delta \ln \frac{w_c^m}{w_c^f} \; = \; \phi_g \sum_j \big(m_{cj} - f_{cj}\big) a_j, \qquad \phi_g > 0, \tag{3}$$

where ϕ_g is the inverse of the sum of the elasticities of relative demand and relative supply between the two. The sign of the right-hand side depends on which industries carry the technology: the male relative wage rises if and only if adopting industries account for a larger share of local male employment than of local female employment. The model therefore imposes no sign on the bias, and the right-hand side of Equation (3) can be computed directly from the same data that build the shifter. That computation shows the average woman working in industries adopting at 1.162 percent against 1.067 percent for the average man, so the between-industry component is slightly negative, and it is negative in 96.9 percent of commuting zones. On this margin alone AI adoption over this window leans a little more toward women than men, because health care and education carry 28.6 percent of the shifter's employment weight. Whatever raises men's relative position must therefore operate within industries, since the industry mix pushes the other way, and a treatment that varies only across industries can sign the incidence without decomposing it. Section 5.3 returns to this. The same construction, repeated across occupational exposure classes, gives a rise in the relative wage of workers in AI-exposed occupations and, because relative supply slopes upward, a rise in their share of local employment.

Households, home time and risk. Individuals have one unit of time and constant absolute risk aversion, $U(c) = -e^{-\gamma c}$ with $\gamma > 0$. An individual of sex g in exposure class $x \in \{H, L\}$ who devotes t of that time to household production earns

$$y = w_{gx}(1 - t) + u, \qquad u \sim \mathcal{N}\big(0, \sigma_x^2\big), \tag{4}$$

where u collects uninsurable labour market risk: displacement, task reassignment and permanent wage shocks. Household production yields κt^{α} with $\kappa > 0$ and $\alpha \in (0, 1)$, and what it yields is enjoyed by everyone in the household, so a couple obtains twice the return a single obtains from the same hour. Married partners share consumption with economies of scale $\theta \in (1, 2]$: each consumes $(y^m + y^f)/\theta$, where $\theta = 1$ is a pure household public good and $\theta = 2$ purely private consumption. Two people in the same exposure class face shocks correlated at $\rho \in (0,1)$, strictly below one because firm-level integration and individual task replacement remain idiosyncratic; across classes the shocks are orthogonal. The diffusion of AI into the exposed class is a change in the pair (w_H, σ_H^2), and nothing in the model dictates the sign of either movement.

Because the earnings shock is additive and preferences are exponential, the allocation of time is independent of risk and every value reduces to a certainty equivalent. A single with wage w chooses $t_s(w) = (\kappa\alpha/w)^\nu$, where $\nu \equiv 1/(1-\alpha)$, and obtains

$$V\big(w, \sigma^2\big) = w\big(1 - t_s(w)\big) + \kappa\, t_s(w)^\alpha - \frac{\gamma}{2}\sigma^2. \tag{5}$$

A couple in which the man is in class k and the woman in class l chooses $t_c(w) = \theta^\nu t_s(w)$ for each spouse and obtains the joint certainty equivalent

$$Y(k,l) = \frac{2}{\theta}\sum_g w^g\big(1 - t_c(w^g)\big) \; + \; 2\kappa \sum_g t_c(w^g)^\alpha \; - \; \frac{\gamma}{\theta^2}\big(\sigma_k^2 + \sigma_l^2 + 2\rho_{kl}\sigma_k\sigma_l\big). \tag{6}$$

Married people put more time into the home than singles do, since $\theta^\nu > 1$ and the return is enjoyed twice, and within a couple the partner with the lower wage puts in more than the partner with the higher wage. Both follow from allocating time efficiently, and neither requires an assumption about who is suited to what.

The marital surplus. Matching is frictionless and utility is transferable, so what governs whether a match forms is its surplus, the difference between the joint value of the union and the two outside options:

$$S(k,l) = Y(k,l) - V\big(w^m, \sigma_k^2\big) - V\big(w^f, \sigma_l^2\big). \tag{7}$$

Substituting Equations (5) and (6) and collecting terms, the surplus splits into a contribution from each spouse's earnings capacity and a term in the two variances,

$$S(k,l) = \underbrace{D\big(w^m\big) + D\big(w^f\big)}_{\text{earnings}} \; + \; \underbrace{\frac{\gamma}{2}\big(\sigma_k^2 + \sigma_l^2\big) - \frac{\gamma}{\theta^2}\big(\sigma_k^2 + \sigma_l^2 + 2\rho_{kl}\sigma_k\sigma_l\big)}_{\text{risk pooling, } R(k,l)}, \tag{8}$$

where the earnings contribution of one spouse is

$$D(w) \;\equiv\; \frac{2}{\theta} w\big(1 - t_c(w)\big) + 2\kappa\, t_c(w)^\alpha - w\big(1 - t_s(w)\big) - \kappa\, t_s(w)^\alpha. \tag{9}$$

By the envelope theorem,

$$D'(w) = \frac{2}{\theta}\big(1 - t_c(w)\big) - \big(1 - t_s(w)\big) = \Big(\frac{2}{\theta} - 1\Big) - \big(2\theta^{\nu-1} - 1\big)\, t_s(w). \tag{10}$$

Set $\kappa = 0$ and the household produces nothing at home: $t_s = t_c = 0$, $D(w) = w(2/\theta - 1)$, and the surplus collapses to the textbook expression in which the two spouses' wages enter identically (Becker, 1973; Chiappori et al., 2017). In that case the incidence of a shock across the sexes is irrelevant, since only the sum of the two wages appears anywhere. Home production breaks the symmetry, and it does so through the wage alone: t_s is decreasing in w, so D' is increasing in it, and the household converts the earnings of the higher-earning spouse into consumption at the higher rate.

The two channels.

Proposition 1. *Let $\theta \in (1, 2)$, $\alpha \in (0, 1)$, $w^m > w^f$, and $w_H \geq w_L$ within each sex.* (i) *D' is strictly increasing, so the surplus responds more to the man's wage than to the woman's,*

$$\frac{\partial S}{\partial w^m} - \frac{\partial S}{\partial w^f} = \big(2\theta^{\nu-1} - 1\big)\big[\, t_s(w^f) - t_s(w^m)\,\big] \;>\; 0,$$

and $\partial S/\partial w^g > 0$ whenever $t_s(w^g) < (2/\theta - 1)/(2\theta^{\nu-1} - 1)$. (ii) *The earnings term is modular in the partners' exposure classes,*

$$\big[D(w^m_H) + D(w^f_H)\big] + \big[D(w^m_L) + D(w^f_L)\big] - \big[D(w^m_H) + D(w^f_L)\big] - \big[D(w^m_L) + D(w^f_H)\big] = 0,$$

so a rise in the exposed class's wages raises the surplus of a match in proportion to the number of exposed partners, by most for H–H and by least for L–L. (iii) *The risk term is strictly submodular in the partners' exposure classes for any positive spousal correlation,*

$$\frac{\partial}{\partial \sigma_H^2}\Big[\, S(H,H) + S(L,L) - S(H,L) - S(L,H)\,\Big] \;=\; -\,\frac{2\gamma\rho}{\theta^2} \;<\; 0,$$

for every admissible θ and every $\rho > 0$.

The proof is in Appendix A. The three parts place very different demands on parameters. Part (i) asks only that marriage deliver some consumption gain and that market work dominate home time at the wage in question. The household literature places the equivalence scale for a two-adult household between 1.3 and 1.6, with structural estimates of indifference scales centring near $\theta \approx 1.41$ (Browning et al., 2013). The size of the wedge turns on the ratio of the two spouses' home time. With home production calibrated so that a married woman spends 1.8 times as long on household work as a married man, which is close to what American time-use data report (Aguiar and Hurst, 2007), and with the female-to-male earnings ratio set at 0.75, a dollar of male earnings moves the surplus by 1.58 times as much as a dollar of female earnings at $\theta = 1.41$. The male derivative stays positive up to $\theta \approx 1.70$ and the female derivative turns negative at $\theta \approx 1.57$, so the independence effect of the earlier literature (Becker, 1981; Greenwood and Guner, 2004) appears here as a region of the parameter space, with no separate assumption required. The left panel of Figure 3 plots both derivatives; the wedge between them measures the additional surplus a male-biased labour-demand shock generates over a sex-neutral one of the same size.

Part (iii) asks for nothing beyond $\rho > 0$, which makes the risk channel testable here. Whether a rise in earnings risk raises or lowers the surplus of a match between two exposed workers turns on a comparison between θ and $\sqrt{2(1+\rho)}$: at $\rho = 0.15$ the channel is positive only above $\theta = 1.52$ and at $\rho = 0$ only above 1.41, so within the range the evidence supports (Hyslop, 2001; Blundell et al., 2016) the sign is unresolved. That comparison governs only the response to a movement in σ_H^2, however, and σ_H^2 has an observable counterpart in the within-commuting-zone dispersion of earnings. The counterpart is imperfect: a cross-section across workers and the variance facing one worker need not move together, and Section 5.3 treats it with that in mind. Part (iii) does not depend on that counterpart. Whatever dispersion does, and wherever θ and ρ lie, a shock working through risk has to place the marriages it makes in couples whose partners are differently exposed. The right panel of Figure 3 translates part (iii) into observed matches. Two shocks scaled to raise total marriages equally place them in very different couples: the earnings channel adds them in proportion to the couple's exposure, while the risk channel concentrates them in mixed pairs and thins both same-type pairings. Because the totals are equal by construction, the composition of matches is the discriminating observable.

The marriage market. Surpluses become marriages in a matching market, one per commuting zone. Men are indexed by their exposure class x and women by theirs, y, with local measures n_{cx} and m_{cy}; the sex ratio $\sum_x n_{cx}/\sum_y m_{cy}$ need not be one. On top of the systematic surplus, each individual draws an idiosyncratic taste for each class of partner and for remaining single, independently and from a type-I extreme value distribution. Choo and Siow (2006) show that equilibrium then has a closed form. Writing μ_{cxy} for the measure of (x, y) marriages and μ_{cx0}, μ_{c0y} for those who marry nobody,

$$\mu_{cxy} = \sqrt{\mu_{cx0}\,\mu_{c0y}}\ \exp\!\left(\tfrac{1}{2}S_{cxy}\right), \qquad \mu_{cx0} + \sum_y \mu_{cxy} = n_{cx}, \qquad \mu_{c0y} + \sum_x \mu_{cxy} = m_{cy}, \quad (11)$$

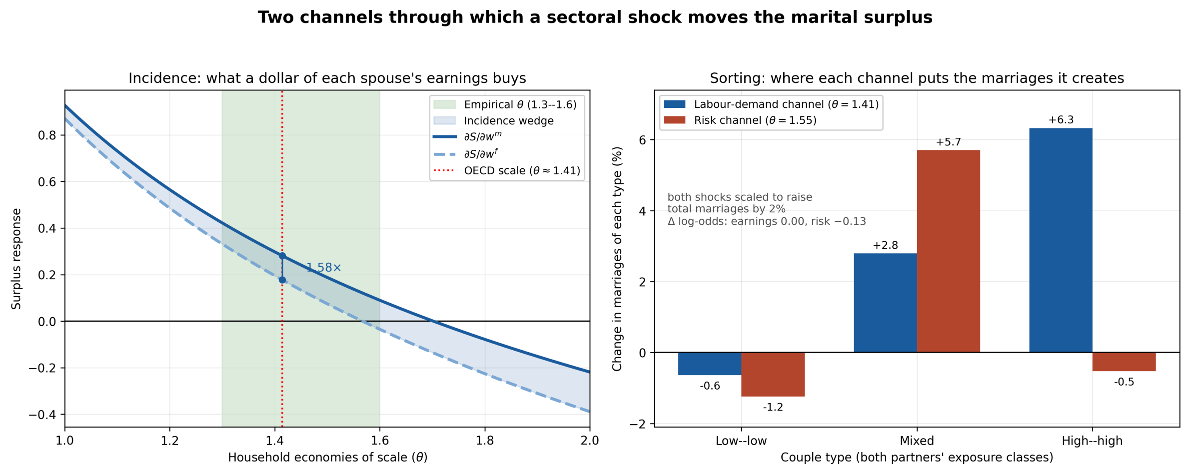

Figure 3: Two channels through which a sectoral shock moves the marital surplus. Left panel: the marginal contribution of each spouse's earnings, $\partial S/\partial w^m$ and $\partial S/\partial w^f$ from Equation (10), against household economies of scale θ. Home production is calibrated at $\alpha = 0.5$ and $\kappa\alpha/w^m = 0.27$ with the female-to-male earnings ratio set at 0.75, which puts married men and women at 0.15 and 0.26 of their time endowment in household work, a ratio of 1.8; the shaded wedge between the curves is the incidence asymmetry, a factor of 1.58 at $\theta = 1.41$. Right panel: the matching market of Equation (11) solved under each channel, with both shocks scaled so that total marriages rise by two percent. The labour-demand shock raises the exposed class's wages at $\theta = 1.41$; the risk shock raises σ_H^2 at $\theta = 1.55$, just above the threshold $\sqrt{2(1+\rho)}$ at which greater risk raises marriage, so the two shocks are observationally equivalent on the total. The earnings channel adds marriages in proportion to the couple's exposure and leaves the log-odds cross-difference of Equation (12) exactly unchanged; the risk channel concentrates them in mixed couples at the expense of both same-type pairings and lowers it. Calibration: a third of each sex in the exposed class, a 25 percent class wage premium, $\rho = 0.15$, and the systematic value of remaining single set so that 55 percent are married at baseline. Shaded band in the left panel marks the bounds on θ from the household literature.

with systematic surplus $S_{cxy} = D(w^m_{cx}) + D(w^f_{cy}) + R(x,y) + \zeta\,\mathbf{1}\{x = y\}$, the last term collecting the non-pecuniary return to similarity that generates the assortative matching the data display (Eika et al., 2019; Chiappori et al., 2018). The adding-up conditions impose market clearing, and the singles constitute the participation margin: nobody is assigned a partner, and the number who go unmatched is determined jointly with who marries whom. Equilibrium transfers clear the market, so the pattern of matches is an equilibrium object, and no matching rule is assumed (Galichon and Salanié, 2022).

Two features of Equation (11) carry the results below. The first is that the degree of assortative matching depends on the surplus alone, because the equilibrium terms cancel from the local log odds,

$$\ln \frac{\mu_{cHH}\,\mu_{cLL}}{\mu_{cHL}\,\mu_{cLH}} = \tfrac{1}{2}\big[\,S_{cHH} + S_{cLL} - S_{cHL} - S_{cLH}\,\big], \tag{12}$$

so parts (ii) and (iii) of the Proposition are statements about an observable. The earnings channel leaves this object unchanged, because it raises the number of matches between exposed workers by raising the value of those workers to any partner, without altering how well suited they are to one another. The risk channel has to lower it.

The second is that the pools of singles govern the response of the market to a shock, however large the population. Writing $M_c = \sum_{xy} \mu_{cxy}$ for the number of marriages and s^m_c, s^f_c for the two pools of singles, a uniform rise dS in the systematic surplus raises marriages by

$$d \ln M_c \;=\; \frac{dS/2}{1 + \frac{1}{2} M_c\big(1/s^m_c + 1/s^f_c\big)} \;>\; 0 \tag{13}$$

and lowers the number of singles of each sex; the expression is exact when each side of the market is homogeneous and holds to first order otherwise. An increase in the measure of men on the market works through the same denominator,

$$\frac{d \ln M_c}{d \ln n^m_c} \;=\; \frac{n^m_c / s^m_c}{2 + M_c\big(1/s^m_c + 1/s^f_c\big)}, \tag{14}$$

which is larger the scarcer single men are. Equation (14) is the formal content of the argument that a shortage of marriageable men depresses marriage (Wilson, 1996; Autor et al., 2019): when men are the short side of the market, adding to their number raises marriages by more than adding to the number of women would. A shock that raises male employment therefore reaches marriage through a margin distinct from the wage, a margin that is relevant here because exposure moves male full-time employment as well as male earnings.

Union stability. Matches formed in the first period meet, in the second, a match-specific quality shock ξ drawn from a distribution G with density g, representing the non-pecuniary value of the partnership. A couple separates when the realised joint value of the marriage falls below that of the outside options,

$$Y(k,l) + \xi < V\left(w^m, \sigma_k^2\right) + V\left(w^f, \sigma_l^2\right), \tag{15}$$

which is $\xi < -S(k,l)$, so a fraction $G(-S)$ of unions dissolve and the ratio of divorced to married individuals is $G(-S)/[1-G(-S)]$. That ratio is strictly decreasing in the surplus,

$$\frac{\partial}{\partial w^m} \frac{G(-S)}{1-G(-S)} = -\frac{g(-S)}{\left[1-G(-S)\right]^2} D'\left(w^m\right) \;<\; 0, \tag{16}$$

by part (i) of the Proposition. Spouses tolerate lower match quality once the economic value of the union has risen, so formation and stability move together and in the same direction as the labour-demand shock. The two marital outcomes reported below are therefore not independent tests of the same hypothesis, and I read them as one.

From the match surplus to the commuting-zone regression. Collecting the three blocks gives the object the empirical work estimates. A national adoption vector $\{a_j\}$ raises the local wage index by ϕB_c through Equation (2) and distributes it across the sexes through Equation (3); the surplus of every match in the zone rises by $D'(w_c^m)\,dw_c^m + D'(w_c^f)\,dw_c^f$; and Equation (13) turns that into marriages. I write λ_c^g for the elasticity of the sex-g wage to exposure implied by Equations (2) and (3), and Ψ_c for the damping factor in Equation (13); to first order

$$\Delta \ln M_c \;=\; \Psi_c \Big[D'\left(w_c^m\right) w_c^m \lambda_c^m + D'\left(w_c^f\right) w_c^f \lambda_c^f \Big] B_c \;\equiv\; \beta_c B_c, \tag{17}$$

which suppresses the employment margin of Equation (14) for legibility, since it enters additively and with the same sign. Equation (17) justifies regressing long differences in local marital outcomes on a shift–share built from national industry adoption and baseline local shares, and it also fixes the coefficient's interpretation: the response of a local marriage market to a local labour-demand shock, damped by how much slack that market has. The only route by which a_j reaches a household is the local labour market, and no term refers to the industry a given worker is employed in, which is why a specification comparing industries within a commuting zone should return nothing, and why the aggregate estimate cannot be read as the sum of own-industry effects. Because Ψ_c and the two elasticities vary across commuting zones, so does β_c; the estimand is a population-weighted average of it, and the baseline demographic controls of Section 4.2 capture the observable part of that heterogeneity.

Hypothesis. The framework yields three implications, and Sections 5 and 5.3 check the sign of each. First, commuting zones with a higher baseline concentration of employment in AI-adopting industries should see a larger rise in marriage and a larger fall in the divorced-to-married ratio, and the rise should be larger the more the local labour-demand expansion favours men. Second, since both channels can raise marriage, the operative channel is identified in the labour market. A shock working through the level of earnings raises employment and men's relative earnings and displaces nobody from the exposed occupations. A shock working through risk must raise measured dispersion, which is the observable counterpart of σ_H^2, and must push employment out of the exposed class. The dispersion test is one-directional: a rise in σ_H^2 implies a rise in dispersion, so flat dispersion would rule the risk channel out, while widening dispersion also arises from a demand expansion concentrated at the top of the distribution and therefore settles nothing. Third, the two channels place the marriages they create in different couples. The earnings term is modular, so a demand shock raises the surplus in proportion to the number of exposed partners and its effect should be concentrated among couples in which both partners work in exposed occupations, while couples in which neither does gain nothing directly and lose ground to the rest of the market. The risk term is submodular for any positive spousal correlation, so a risk shock has to do the opposite and concentrate in mixed couples. Nothing in that last comparison turns on the sign of the risk channel or on where θ and ρ lie, which makes it the more informative of the two tests.

4 Data & Measurement

The observational units are the 741 U.S. commuting zones (1990 definitions). I study long differences over 2010–2021 for three reasons. First, the Annual Business Survey (ABS) Technology Module provides a harmonised snapshot of AI adoption by industry for reference year 2020 for the United States, and the Eurostat ICT Usage in Enterprises survey provides a comparable 2021 module for Europe. Second, the beginning of the window coincides with the onset of modern AI diffusion documented in Section 2. Third, a long-difference design filters out business-cycle noise and attenuates concerns about short-run feedback between local conditions and adoption, while still capturing the decade-scale shift associated with AI.

Data from IPUMS/ACS. I build family and labour market outcomes from ACS microdata (IPUMS) (Ruggles et al., 2025), using person weights and restricting to ages 18–64, and map individuals to commuting zones via the PUMA-to-CZ crosswalk; when a PUMA spans multiple CZs, person weights are allocated proportionally to crosswalk shares so that composition reflects the underlying geography.

The marital outcomes are the share of adults who are married, the share in a first marriage, and the ratio of the divorced share to the married share. The labour market outcomes are the employment rate (worked in the previous year), the full-time rate (usually working 35 hours or more), the mean of log annual wage and salary income among those with positive earnings, the within-commuting-zone standard deviation of that variable, the share of adults with zero earnings, and the share of employment in occupations above the 66^{th} percentile of the Webb (2019) AI exposure distribution. Wage variables are constructed over earners aged 18–64 with valid, strictly positive wage and salary income. Appendix Table B1 reports summary statistics for the change in each outcome over 2010–2021.

Baseline controls summarise the demographic and educational structure of each commuting zone in the base year of the window being estimated: the female share, the share aged 60 and over, the BA+ share, and the male and female shares of employment in high-exposure occupations. Measuring controls at each window's own base year means the pre-period placebo of Section 5.2 conditions on nothing realised after that window begins. Descriptive statistics are reported in Appendix Table B2.

Industry AI adoption. I take industry-level AI adoption for the United States from the ABS Technology Module for reference year 2020. The survey asks firms about their active use of five specific technological classes: machine learning, machine vision, natural language processing, voice recognition and automated guided vehicles. Within each technology, the industry adoption rate is firm-weighted: adopting firms are summed across the ABS rows composing the industry and divided by the summed firm count, so that nested child sectors are never double counted. The industry shifter $\text{AIAdopt}_j^{\text{US}}$ is then the unweighted mean across the five technologies of those rates. For arts, entertainment and recreation, where the five rates are 0.2, 0.4, 0.4, 0.3 and 1.1 percent, the shifter is their mean, 0.48 percent.

Two features of this construction bear on interpretation. The shifter is an intensity measure, the average rate of adoption across five non-mutually-exclusive technologies. The share of firms using at least one technology would be the more natural object, but the ABS publishes the five technologies separately and not their union, so that share is not observed; the maximum across the five is a lower bound on it and their sum is an upper bound. Appendix Table B3 re-estimates the main specification on both bounds, on each technology taken separately, and on the three cognitive technologies alone, and every variant delivers the same conclusion.

The measure covers 46 industries. Appendix Table B4 lists all of them together with the adoption rate of each technology and the resulting shifter. Figure 4 displays the same information collapsed to broad NAICS sectors for legibility; the estimation uses the 46-industry version throughout. The matrix structure shows that different industries integrate different subsets of AI capabilities: information and professional services rely heavily on machine learning and natural language processing, while manufacturing and transportation rely on machine vision and automated guided vehicles. This sorting of specific tools to specific sectoral needs indicates that the module captures targeted, task-replacing technology and not a uniform increase in general computing capital. The point bears on interpretation, since the industries that score highest are not exclusively the cognitive-services industries.

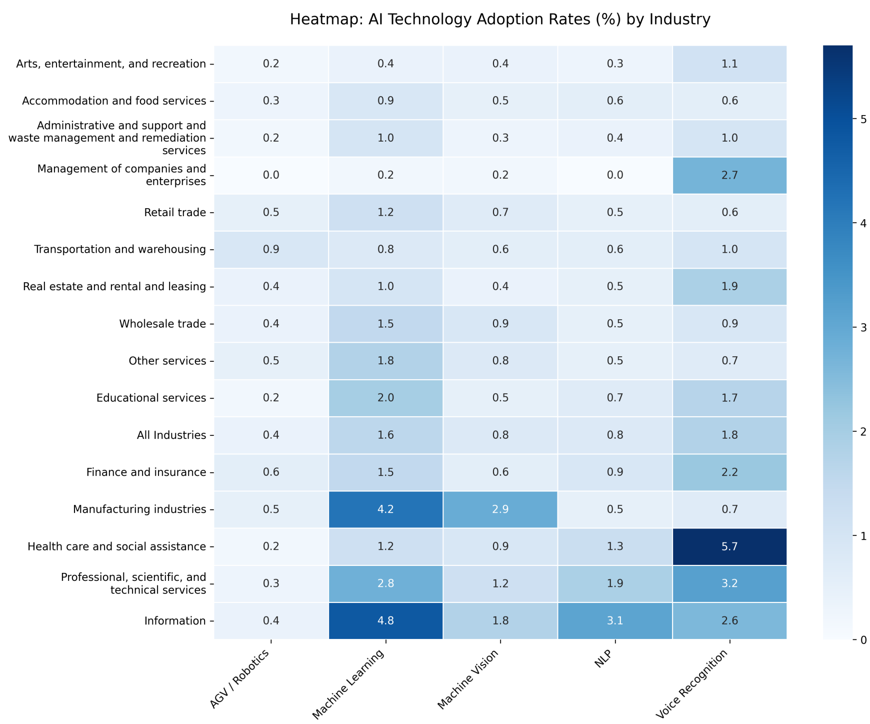

Figure 4: Heatmap of AI adoption rates by technology class and industry sector. Values represent the percentage of firms in the sector reporting active use. Broad NAICS sectors are shown for legibility; the 46 industry cells used in estimation, with the same decomposition, are in Appendix Table B4. Data source: ABS Technology Module.

Local industry shares. Local industry composition is measured from County Business Patterns (CBP). For each commuting zone c and industry j, I compute baseline employment shares $s_{cj,2010} = L_{cj,2010} / \sum_{j'} L_{cj',2010}$.

4.1 Shift–share AI exposure

The regressor is a Bartik-style exposure that weights national industry adoption by local pre-period industry structure:

$$\text{AI Exposure}_c = \sum_{j=1}^{46} s_{cj,2010} \cdot \text{AIAdopt}_j^{\text{US}}.$$

A commuting zone scores high if it was initially specialised in industries that subsequently adopted AI intensively. Using 2010 shares addresses reverse causality, since shares are fixed before the adoption window, and reduces measurement error relative to year-by-year local adoption series. A one-unit change in exposure is the difference between a commuting zone with no exposure and one in which every firm in every industry uses all five technologies; since observed exposure runs from zero to 0.0137 with a standard deviation of 0.0023 (Table 1), magnitudes are reported throughout by scaling coefficients to a one-standard-deviation increase.

Figure 5 plots commuting-zone exposure alongside the long-run changes in marriage shares and divorce-to-marriage ratios. Exposure is concentrated in the industrial Midwest, Texas and parts of the Southeast, and is comparatively lower in coastal regions and the Mountain West, a geography that follows from the sectoral pattern of Figure 4.

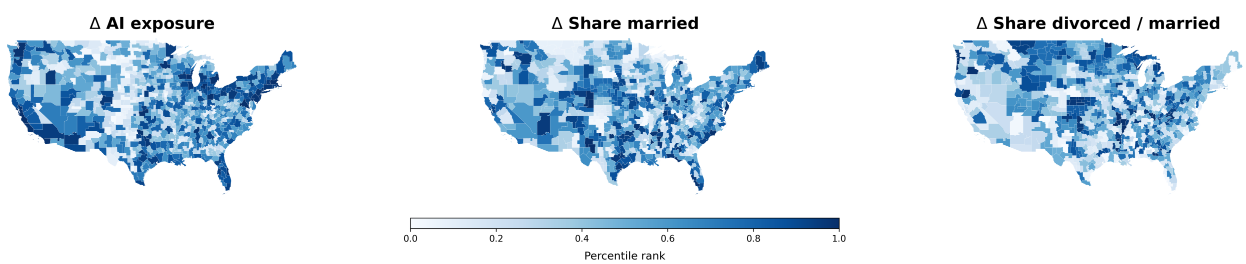

Figure 5: **AI exposure, marriage and divorce across commuting zones.** The three panels map commuting-zone AI exposure and the 2010–2021 changes in the share married and in the divorced-to-married ratio, each shaded by percentile rank. Exposure is the shift-share defined in Section 4.1. Sources: ABS Technology Module, County Business Patterns and ACS microdata.

Table 1: Summary statistics: AI shift–share exposures (CZ level, 2010–2021)

Variable	N	Mean	Std. dev.	Min	Max
AI Exposure	741	0.0074	0.0023	0.0000	0.0137
AI Exposure^{IV}	741	0.0363	0.0102	0.0000	0.0733

Endogeneity remains a concern if industries expand AI in response to unobserved U.S.

factors that also shift family outcomes. To address this, I instrument AI Exposure$_c$ with a foreign analogue constructed from European industry adoption:

$$\text{AI Exposure}_c^{IV} = \sum_j s_{cj,2010}^{\text{NACE}} \cdot \text{AIAdopt}_j^{\text{EU}}.$$

The European measure is taken from Eurostat's ICT Usage in Enterprises survey for 2021 and is the share of enterprises with ten or more employees reporting the use of at least one AI technology, for the EU-27 aggregate. Two details of the construction matter for interpretation. First, Eurostat does not publish the sectoral breakdown by individual technology, so the European measure is an any-tool share and is therefore not built in the same way as the U.S. shifter. Second, Eurostat's sectoral detail is coarser, covering 27 NACE cells against 46 U.S. industries, so the instrument is formed on a separate share matrix aggregated to the matched cells. This is why the two shift-shares differ in level, averaging 0.0074 and 0.0363 respectively. Appendix Table B5 reports the full NACE-to-NAICS correspondence and the European adoption rate for each cell; the country weighting is Eurostat's own EU-27 aggregate and no further weighting is applied.

Figure 6 plots the endogenous exposure against the instrument at the commuting-zone level. The population-weighted first stage has a slope of 0.146 with an R^2 of 0.70 and a first-stage F statistic of 334.0.

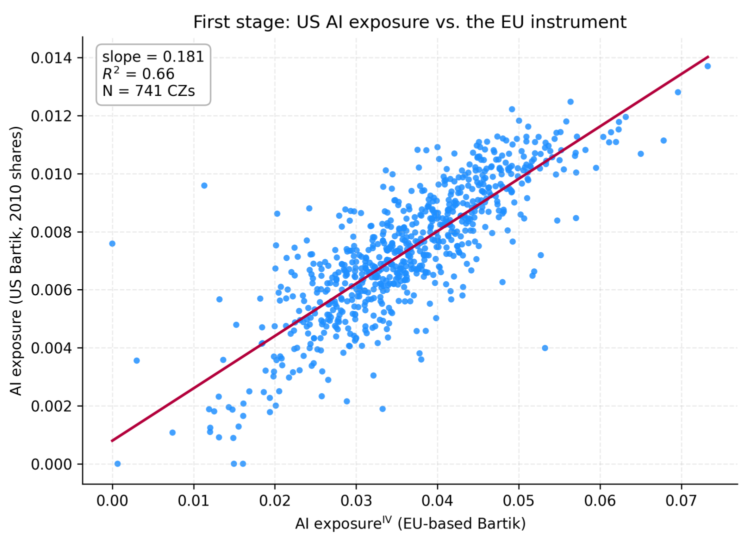

Figure 6: Relationship between the instrumented exposure (AI Exposure^{IV}) and the endogenous exposure (AI Exposure). Points are commuting zones; the annotation reports the OLS slope and R^2 from a linear regression.

4.2 Specification

I estimate the long-run relationship between AI exposure and outcomes in long differences:

$$\Delta_b^t Y_c \;=\; \beta\,\mathrm{AIExposure}_c \;+\; \gamma' X_c^b \;+\; \varepsilon_c,$$

where Y_c is a commuting zone's outcome and X_c^b collects predetermined base-year controls (female, 60+, and BA+ shares; baseline high-AI occupational shares by sex). Using long differences sweeps out CZ-invariant unobservables, and fixed base-year controls absorb baseline compositional differences across places. I weight regressions by base-year commuting-zone population and cluster standard errors by state. Endogeneity in $\mathrm{AIExposure}_c$ is addressed with two-stage least squares using the EU-based shift-share, so that the first stage projects U.S. exposure on its European analogue and the second stage relates outcome changes to the predicted component.

What the instrument delivers. The logic of the European instrument is that global, supply-side technological improvements generate common cross-industry patterns of adoption across advanced economies, so that the instrument isolates the component of U.S. adoption that is orthogonal to U.S. commuting-zone shocks. That logic is sound as far as it goes, but its reach is narrower than in the robotics case, where some European countries were technological leaders relative to the United States and movements in European adoption could plausibly be treated as shocks independent of domestic conditions (Acemoglu and Restrepo, 2020). In AI the United States is the leader, and the industries that adopt intensively in Europe are largely the ones that do so in the United States, for the same underlying reasons: the technology is more valuable where tasks are codifiable and data are abundant. The European measure is therefore best understood as a second, independently collected measurement of much the same industry ranking, supplying little variation orthogonal to that ranking. It corrects classical measurement error in a single-survey shifter, which is why the 2SLS estimates exceed the OLS estimates throughout.

The pattern of attenuation is itself informative: adding the full control set moves the OLS estimate for the share married from 4.592 to 1.718, a fall of 63 percent, against 5.572 to 3.786 for 2SLS, a fall of 32 percent. If the controls were absorbing true treatment variation, both estimators would attenuate alike; the larger OLS fall is expected when part of what the controls remove from the OLS estimate is classical error in the shifter, which the instrument already purges. The instrument also guards against the narrower concern that U.S. industry adoption responded to U.S. local demand conditions in the same years in which local marriage rates moved. It does not, on its own, address the broader concern that industries adopting

AI differ systematically from those that do not.

What the setting delivers. Four features of the setting work in the design's favour. First, baseline shares are fixed in 2010, before diffusion began, so the local industrial structure that carries the shock cannot have adjusted to it. Second, adoption is measured directly: the ABS asks firms about the active use of named technologies, whereas an exposure index infers vulnerability from task descriptions. Third, the instrument supplies a second, independently collected measurement of the industry ranking, which purges classical measurement error from the shifter. Fourth, the window has a counterfactual on each side: the identical exposure predicts nothing over 2000–2007, and the estimates build from 2012 onward, well before the pandemic.

What the design identifies. The exogenous-shifts framework (Borusyak et al., 2023; Adão et al., 2019) asks that industry shocks be as good as randomly assigned conditional on controls, and that there be enough of them. Neither condition is comfortable here. Appendix B.1 reports the effective number of shocks, the Rotemberg weights, a shock-permutation test and two ways of conditioning the shifter on industry structure. Those diagnostics imply that the estimand is the effect of specialising in the industries that went on to adopt AI, relative to other commuting zones over the same period; it carries no claim about the effect of AI holding constant everything correlated with adoption.

Interpretation of magnitudes. Two points of interpretation follow from the design. First, every estimate below is differential. A shift-share design identifies how outcomes differ across commuting zones by exposure and cannot recover the average effect of AI adoption on the country as a whole, since that common component is absorbed. The results are consistent with AI reducing marriage everywhere while reducing it least in the most exposed places, and no statement about the share of a national trend offset by AI can be supported by this design. Second, sample-mean changes are reported alongside the coefficients to give a sense of scale; they do not license an absolute interpretation. Inference is reported three ways: clustered by state throughout, with the Adão et al. (2019) exposure-robust correction in Appendix Table B6, and by geographic permutation following Goldsmith-Pinkham et al. (2020) in Figure 9.

5 Results

This section presents the empirical findings. It begins with the local marriage market, relating commuting-zone AI exposure to changes in marriage and divorce, and subjects those estimates to a suite of robustness checks. The local labour market follows: the outcomes that identify the operative channel, the tests that discriminate between the two channels of Section 3, and the within-commuting-zone design that asks whether the effect operates at the level of the local market or of the individual worker.

5.1 Local marriage market results

Table 2 reports long-difference estimates relating local exposure to AI adoption to changes in marital outcomes between 2010 and 2021. Each column presents a commuting-zone regression weighted by base-year population, with standard errors clustered by state. Panel A reports OLS estimates and Panel B the corresponding 2SLS estimates, instrumenting U.S. exposure with its European analogue. Columns (1)–(3) analyse the change in the overall share of married adults, columns (4)–(6) the share in a first marriage, isolating entry into legal marriage and excluding remarriages, and columns (7)–(9) the ratio of divorced to married individuals.

In the specifications with full demographic and occupational controls, the estimates indicate a positive relationship between AI exposure and marriage and a negative relationship with divorce prevalence. For the share in a first marriage (column 6), the OLS coefficient is 2.087 and the 2SLS estimate is 5.758, significant at the 1 percent level. Scaling by one standard deviation of exposure implies a 1.30 percentage point relative increase, against a sample-mean change of -1.18 points over the same period. For the divorced-to-married ratio (column 9) the 2SLS estimate is -8.262, a -1.87 point relative decline against a mean change of -0.49 points.

The overall married share (column 3) rises by 3.786, significant at 5 percent but the weakest of the three results. The ordering is expected, because the married stock adjusts slowly and pools long-standing marriages with new ones, whereas the first-marriage share isolates the entry margin on which a labour-demand shock should operate most directly. The pattern that the effect on the stock is smaller and less precisely estimated than the effect on the flow recurs throughout the robustness checks below.

Table 2: AI exposure and changes in marital outcomes, 2010–2021

	Δ Share married			**Δ Share 1^{st} married**			**Δ Share divorced / married**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	4.592***	3.789***	1.718*	3.542***	3.738***	2.087**	-5.171***	-5.348***	-4.888***
	(0.440)	(0.613)	(0.908)	(0.563)	(0.694)	(0.979)	(1.010)	(1.306)	(1.466)
Female share		0.132	-0.0758		-0.0102	-0.0803		0.0417	0.103
		(0.137)	(0.105)		(0.138)	(0.106)		(0.189)	(0.160)
Age 60+ share		-0.154	-0.0918		0.0504	0.0697		-0.0268	-0.0279
		(0.146)	(0.153)		(0.135)	(0.142)		(0.210)	(0.215)
BA+ share			0.128***			0.0443			-0.0155
			(0.0355)			(0.0378)			(0.0615)
High-AI occ., women			-0.0516			0.0809			-0.0835
			(0.0578)			(0.0667)			(0.122)
High-AI occ., men			-0.0938			-0.0450			0.0735
			(0.0568)			(0.0584)			(0.0951)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.119	0.130	0.187	0.075	0.076	0.098	0.078	0.078	0.082
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	5.572***	4.980***	3.786**	5.047***	5.962***	5.758***	-6.369***	-7.072***	-8.262***
	(0.658)	(0.892)	(1.720)	(0.586)	(0.921)	(1.729)	(1.085)	(1.767)	(2.695)
Female share		0.0785	-0.111		-0.111	-0.143		0.120	0.161
		(0.150)	(0.116)		(0.161)	(0.119)		(0.211)	(0.154)
Age 60+ share		-0.0983	-0.0445		0.155	0.154		-0.108	-0.105
		(0.147)	(0.150)		(0.133)	(0.141)		(0.220)	(0.215)
BA+ share			0.106**			0.00454			0.0210
			(0.0430)			(0.0470)			(0.0755)
High-AI occ., women			-0.0636			0.0596			-0.0639
			(0.0563)			(0.0645)			(0.111)
High-AI occ., men			-0.0941*			-0.0454			0.0739
			(0.0547)			(0.0551)			(0.0925)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.114	0.124	0.176	0.062	0.054	0.061	0.074	0.072	0.067

CZ-level long differences, 2010–2021. Panel A reports OLS estimates and Panel B reports 2SLS estimates instrumenting U.S. exposure with its European analogue. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$. Columns (1), (4) and (7) include no controls; (2), (5) and (8) add the baseline female and age 60+ shares; (3), (6) and (9) add the baseline BA+ share and the sex-specific baseline shares of employment in AI-exposed occupations.

5.2 Robustness checks

This section assesses the principal threats to validity for the marital estimates: pre-existing differences between exposed and unexposed commuting zones, confounding from the Great Recession and from the COVID-19 pandemic, confounding from contemporaneous robotic automation, and inference in a shift-share setting in which shocks occur at the industry level. The labour-market outcomes face the same threats, and Section 5.3 subjects them to the same checks alongside the evidence they support.

Pre-period placebo and the timing of the effect. Table 3 runs the same specification over four windows. The 2000–2007 column is a placebo: AI diffusion had not begun, so exposure should predict nothing, and it does not. Both coefficients are small relative to their standard errors, -0.107 for the share married and -2.154 for the divorced-to-married ratio. The first-marriage share cannot be constructed for this window because the ACS marital-history questions begin only in 2008.

The 2007–2010 column identifies a confounder. The same commuting zones were hit unusually hard by the Great Recession (the labour-market declines are documented in Table 7), so part of what follows the trough is recovery; the recession controls below address this directly.

Table 3: AI exposure and marital outcomes across four windows: pre-recession placebo, recession, pre-COVID and full sample

	Placebo	Recession	Treatment window	
	2000–2007	2007–2010	2010–2019	2010–2021
Δ Share married	-0.107	0.936	0.586	3.786**
	(1.029)	(1.213)	(1.355)	(1.720)
Δ Share 1^{st} married		-4.722	2.468**	5.758***
		(3.558)	(1.240)	(1.729)
Δ Share divorced / married	-2.154	-2.695	-4.705*	-8.262***
	(2.315)	(1.753)	(2.471)	(2.695)

Each cell is a separate 2SLS long-difference regression of the row outcome on AI exposure, with the full control set measured at the base year of that window, so the 2000–2007 column conditions on nothing realised after 2000. The first-marriage share is unavailable before 2008 because the ACS marital-history questions begin then. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Accounting for the Covid-19 shock. A potential confounder in the 2021 endline is the COVID-19 pandemic, which generated idiosyncratic disruptions to household formation, divorce proceedings and regional mobility. The 2010–2019 column of Table 3 restricts the dependent variable to the pre-pandemic long difference. First marriage (2.468) and the divorced-to-married ratio (-4.705) survive; only the overall married share does not (0.586), consistent with the stock-versus-flow ordering noted above.

Figure 7 traces the coefficients by endpoint year, holding the 2010 base fixed. The first-marriage and divorce effects appear from the middle of the decade and are visible well before the pandemic, and the 2020–21 endpoints strengthen effects that were already present; the married share moves latest and least, as a stock should. The traces are noisier than their labour-market counterparts in Section 5.3, which build smoothly from 2012, but they support the same conclusion: the pandemic, whatever it did to these commuting zones, does not generate the estimates.

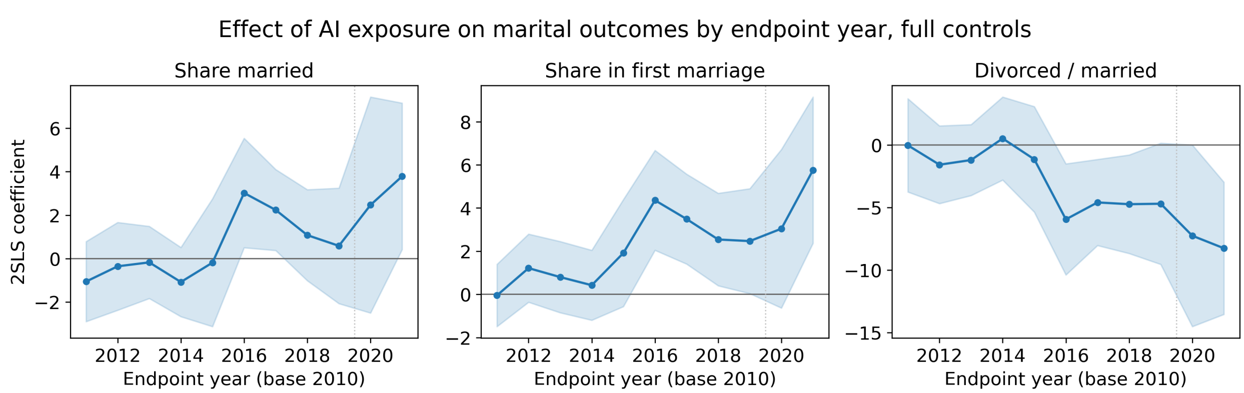

Figure 7: **Effect of AI exposure on marital outcomes by endpoint year.** Each point is the fully controlled 2SLS coefficient from a separate regression of the 2010-to-t long difference on AI exposure; shaded bands are 95 percent confidence intervals from state-clustered standard errors. The dotted vertical line marks the last pre-pandemic endpoint.

A related concern is that because AI-exposed occupations are disproportionately screen-based, exposure might proxy for the local capacity to work from home during 2020–21. Following the classification of Dingel and Neiman (2020), I construct a commuting-zone index of telework capacity as the working-age-population-weighted mean teleworkability of local occupations, and include it as a control. Appendix Table B7 shows that the main estimates are unaffected, as Figure 7 would lead one to expect.

Functional form. The specification is linear in exposure, which imposes the same effect per unit throughout the distribution. Figure 8 relaxes this, replacing the linear term with indicators for population-weighted quintiles of exposure and plotting the resulting bin effects against the fitted line. The marital outcomes are noisy at bin level and their confidence intervals are correspondingly wide, but the linear fit lies inside them at every quintile and the ordering is the expected one. The labour-market version of this check in Section 5.3 is sharper, with a gradient monotone across every quintile. I retain the linear coefficient throughout for comparability, with this figure as the check on it.

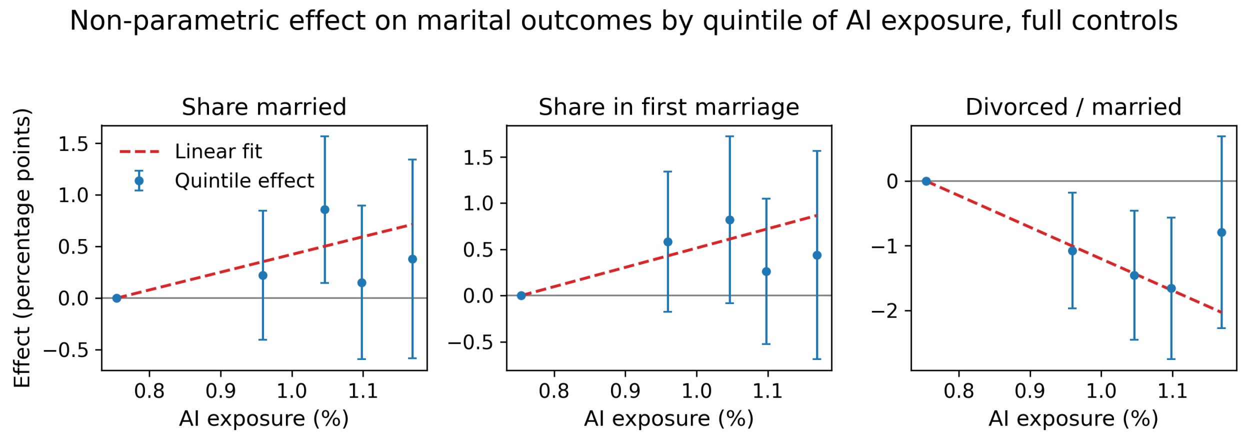

Figure 8: **Non-parametric effect of AI exposure on marital outcomes.** Points are the estimated effects of population-weighted quintiles of AI exposure relative to the lowest quintile, from OLS regressions with the full control set, plotted at the population-weighted mean exposure of each quintile; bars are 95 percent confidence intervals from state-clustered standard errors. The dashed line is the linear fit from the corresponding single-coefficient regression. Effects are in percentage points.

Recovery from the Great Recession. Table 4 addresses the confounder identified above by adding each commuting zone's own 2007–2010 changes in employment, male full-time work and male earnings as controls, so that the comparison is between places that experienced the recession similarly.[1] The estimates attenuate and survive. First marriage falls from 5.758 to 4.416 and the divorced-to-married ratio from -8.262 to -6.372, both still significant at the 1 percent level; the overall married share falls to 2.620 and is no longer distinguishable from zero at the 5 percent level. Recovery from the recession therefore explains part of the effect without eliminating it, and Section 5.3 reports the labour-market counterpart of this check.

[1]The controls are male-specific because the recession differential was largest on the male margins (male full-time work fell by 11.132 against 5.995 for women, male earnings by 18.064 against 14.205), and recovery on those margins could masquerade as the incidence the results document. The choice is also the conservative one: with the overall full-time and earnings changes as controls instead, the marital estimates are essentially unchanged and the earnings gap of Table 8 survives at 15.401 against 11.778.

Table 4: Robustness: marital outcomes controlling for 2007–2010 recession severity, 2010–2021

	Δ Share married			**Δ Share 1^{st} married**			**Δ Share divorced / married**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	1.718*	1.616*	1.052	2.087**	1.895**	1.294	-4.888***	-4.713***	-3.873***
	(0.908)	(0.836)	(0.789)	(0.979)	(0.895)	(0.873)	(1.466)	(1.384)	(1.386)
Δ_{07}^{10} Employment rate		-0.0432	0.0602		-0.0812	0.0392		0.0739	-0.129
		(0.0868)	(0.0803)		(0.0807)	(0.0843)		(0.137)	(0.131)
Δ_{07}^{10} Full-time rate, men			-0.0429			-0.0571			0.119
			(0.0521)			(0.0570)			(0.0865)
Δ_{07}^{10} Mean ln earnings, men			-0.0495***			-0.0500***			0.0609**
			(0.0169)			(0.0173)			(0.0292)
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	3.786**	3.614**	2.620*	5.758***	5.418***	4.416***	-8.262***	-7.955***	-6.372***
	(1.720)	(1.555)	(1.511)	(1.729)	(1.549)	(1.609)	(2.695)	(2.484)	(2.246)
Δ_{07}^{10} Employment rate		-0.0238	0.0660		-0.0470	0.0507		0.0424	-0.138
		(0.0793)	(0.0774)		(0.0696)	(0.0820)		(0.120)	(0.126)
Δ_{07}^{10} Full-time rate, men			-0.0384			-0.0483			0.112
			(0.0528)			(0.0581)			(0.0864)
Δ_{07}^{10} Mean ln earnings, men			-0.0459***			-0.0429***			0.0551**
			(0.0160)			(0.0161)			(0.0267)

Commuting zones specialised in AI-adopting industries lost employment and male earnings sharply in 2007–2010 (Table 7), so part of the 2010–2021 change could be recovery. Columns add the commuting zone's own 2007–2010 changes as controls. Within each outcome, column 1 is the baseline, column 2 adds the employment change and column 3 adds the male full-time and male earnings changes. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Pre-trends. Appendix Table B8 controls for the pre-period change in the respective outcome variable. For the share married and the divorced-to-married ratio the pre-period is 2005–2010; for the share in a first marriage it is 2008–2010, because the ACS marital-history questions used to identify first marriages were introduced only in the 2008 survey year, so 2008 is the earliest year for which a first-marriage share can be measured on a basis consistent with the 2010 and 2021 endpoints. The pre-period change enters strongly and negatively in every column, confirming mean reversion in local marital dynamics, while the AI coefficients are essentially unchanged relative to the baseline.

Robot exposure. A natural concern is that AI adoption co-moves with other capital-intensive technologies. To separate AI from generic automation, Appendix Table B9 controls for the change in local robot exposure between 2010 and 2019, constructed following

Acemoglu and Restrepo (2020). Including it leaves the main AI coefficients essentially unchanged. The coefficient on robot exposure itself is small and, where significant, positive. This is not in tension with Anelli et al. (2024), whose identifying variation is the earlier robot wave in manufacturing-intensive areas; here it is a 2010–2019 change entered conditionally on AI exposure, in a decade in which robot adoption was slower and concentrated in the same places as AI adoption. The control establishes that the AI estimate is not a relabelled robot estimate.

Cost of living and economic trajectories. Because exposure is correlated with local economic conditions, a reasonable concern is that the marriage response reflects cost-of-living dynamics, with individuals formalising unions to share housing costs where those costs are high. This is better understood as a competing mechanism than as a nuisance, and the relevant evidence is Appendix Table B10, which controls for baseline median household income in 2000 and the annualised income growth rate between 2000 and 2010 from Chetty et al. (2014). The estimates are unchanged, which weighs against a pure cost-sharing account.

Predetermined industry shares (2003 baseline). If early AI adoption or anticipatory investment induced sectoral reallocation just prior to 2010, measuring employment shares in that year could introduce endogeneity into the Bartik weights. Appendix Table B11 reconstructs the exposure measure using 2003 industry employment shares while maintaining the 2010–2021 outcome window. The results are stable, confirming that the estimates are driven by persistent historical industrial specialisation rather than by short-term labour market sorting immediately preceding the AI boom.

An occupation-based exposure measure. The treatment measures the AI adoption of the industries a commuting zone was specialised in. An alternative, and the measure an occupational-task literature would reach for, is how AI-exposed the local workforce already was: the 2010 employment-weighted mean of the Webb (2019) occupational percentile. The two capture different things (a shock projected onto local structure against a baseline characteristic of the workforce) and across commuting zones they are close to orthogonal, correlating -0.198. Appendix Table B12 reports the occupational measure entered alone and alongside the shift-share. On its own it explains essentially nothing. Entered jointly with the shift-share, it leaves the shift-share estimates unchanged: first marriage 5.416, the divorced-to-married ratio -7.917, employment 9.528 and the earnings gap 19.844. The results are therefore not a relabelled measure of local occupational exposure.

Industry-shock inference. Because the exposure measure is constructed from a finite set of national industry shocks, standard errors clustered by geography may understate uncertainty if regression residuals are correlated across regions with similar industrial compositions (Adão et al., 2019). Appendix Table B6 reports standard errors calculated using the Adão–Kolesár–Morales procedure, which re-centres inference on the industry-level shocks rather than the regional units. The standard errors increase relative to the state-clustered baseline, as expected under more conservative industry-level inference, and the main coefficients remain statistically significant.

Location-based randomization inference. As a further check on inference that respects the data structure, I implement a location-based randomization test following Goldsmith-Pinkham et al. (2020). This procedure holds the national technological shocks fixed and randomises the geographic assignment of baseline industrial structures. Letting s_{cj} denote the employment share of industry j in commuting zone c and z_j the national shock, the true exposure is $B_c = \sum_j s_{cj} z_j$; I generate a random permutation $\pi(c)$ of commuting zones and construct placebo exposures $B_c^{\pi} = \sum_j s_{\pi(c)j} z_j$. This preserves the joint structure of U.S. and EU adoption, and thus the first-stage relationship, while testing whether the observed responses are driven by the actual geographic alignment of AI-exposed industries. Figure 9 displays the results over 1,000 permutations. The exact permutation p-values are 0.021 for the first-marriage share and 0.016 for the divorced-to-married ratio, with the overall married share marginal at 0.067, the same ordering across outcomes that the state-clustered estimates show. This test addresses sampling error given the treatment; it is distinct from, and does not substitute for, the shock-permutation evidence of Section 4.2, which concerns what the treatment measures.

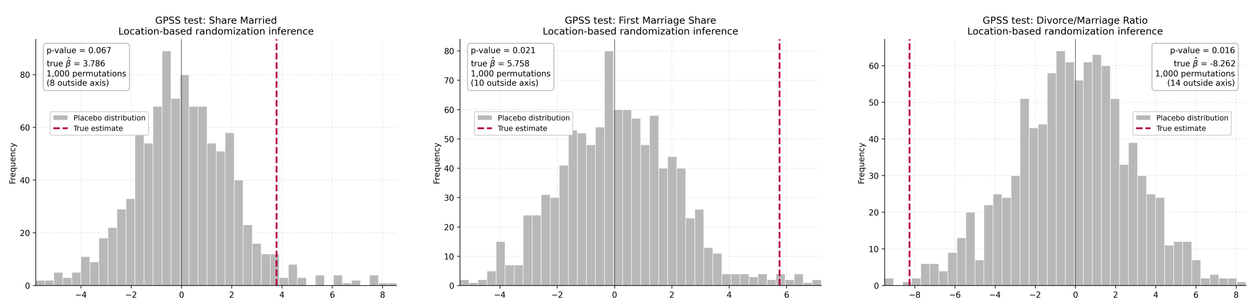

(a) Total Marriage Share (b) First Marriage Share (c) Divorce/Marriage Ratio

Figure 9: **Geographic Location Permutations.** The panels display the empirical distribution of 1,000 placebo 2SLS coefficients generated by randomly reassigning baseline industry shares across commuting zones. The red dashed lines indicate the true 2SLS point estimates. Exact p-values are calculated as the share of placebo coefficients more extreme than the true estimate.

5.3 Local labour market results and mechanism

Section 3 identified two channels through which a sectoral shock can raise the marital surplus, and showed that they carry different empirical signatures. The labour-demand channel requires exposure to raise employment and men's relative earnings. The risk-sharing channel imposes three requirements of its own: exposure must raise the dispersion of earnings, must displace workers out of exposed occupations, and must place the marriages it creates in couples whose partners are differently exposed. Only the last two of those discriminate, for the reason given in Section 3: a demand expansion widens the earnings distribution as readily as a rise in risk does, so dispersion cannot separate them, while displacement and the placement of the new marriages can. This section establishes the labour-demand evidence first, tests the risk channel's requirements against it, and closes by asking whether the effect operates at the level of the individual worker or of the local labour market as a whole.

Local labour demand. Table 5 reports what exposure did to the local labour market. AI exposure raises the employment rate of the 18–64 population by 9.237, the male full-time employment rate by 13.503, and the male–female gap in full-time employment by 8.055, each significant at the 1 percent level. Scaled to a one-standard-deviation increase in exposure, the employment effect is 2.09 percentage points.

Table 5: AI exposure and local labour demand, 2010–2021

	Δ Employment rate			**Δ Full-time rate, men**			**Δ Full-time gap (M−F)**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	2.507**	2.466*	3.681**	-0.0305	1.855	4.117*	-1.340	-0.648	1.862
	(0.997)	(1.359)	(1.565)	(1.767)	(1.921)	(2.356)	(1.285)	(1.517)	(1.789)
Female share		0.345**	0.300*		-0.0415	0.231		-0.0444	0.214*
		(0.167)	(0.177)		(0.354)	(0.205)		(0.215)	(0.120)
Age 60+ share		0.186	0.154		0.517**	0.357*		0.173	0.0584
		(0.140)	(0.135)		(0.198)	(0.191)		(0.151)	(0.142)
BA+ share			-0.0354			-0.264**			-0.205***
			(0.0319)			(0.111)			(0.0697)
High-AI occ., women			0.134			0.616***			0.310***
			(0.0831)			(0.144)			(0.109)
High-AI occ., men			-0.130*			-0.0938			0.0152
			(0.0667)			(0.140)			(0.0845)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.035	0.075	0.103	0.000	0.035	0.132	0.005	0.010	0.074
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	3.774***	4.839**	9.237***	2.673	6.632**	13.50***	0.598	2.398	8.055***
	(1.428)	(1.934)	(2.309)	(2.143)	(3.307)	(4.484)	(1.347)	(2.092)	(2.748)
Female share		0.237	0.204		-0.258	0.0696		-0.182	0.107
		(0.151)	(0.162)		(0.438)	(0.209)		(0.266)	(0.111)
Age 60+ share		0.298**	0.281**		0.742***	0.572***		0.317*	0.200
		(0.142)	(0.142)		(0.238)	(0.215)		(0.174)	(0.154)
BA+ share			-0.0956**			-0.365***			-0.272***
			(0.0389)			(0.141)			(0.0851)
High-AI occ., women			0.102			0.561***			0.274***
			(0.0845)			(0.132)			(0.0994)
High-AI occ., men			-0.130**			-0.0950			0.0144
			(0.0634)			(0.134)			(0.0829)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.026	0.052	0.025	-0.016	-0.002	0.043	-0.006	-0.009	0.024

CZ-level long differences, 2010–2021. The employment rate is the share of 18–64 year olds who worked in the previous year; the full-time rate is the share usually working 35 hours or more. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$. Columns (1), (4) and (7) include no controls; (2), (5) and (8) add the baseline female and age 60+ shares; (3), (6) and (9) add the baseline BA+ share and the sex-specific baseline shares of employment in AI-exposed occupations.

Table 6 shows the same pattern in earnings. Men's mean log earnings rise by 15.422; women's are flat and imprecisely estimated at -4.240 (s.e. 4.040); and the gap between them widens by 19.662, significant at the 1 percent level. The outcome is annual wage and salary income, so it moves with hours as well as with the wage; the incidence-by-sex evidence below recomputes it as a wage per hour, on which close to two thirds of the gap remains. Employment and earnings move in the same direction, which marks the shock as a shift in labour demand; a supply shift would move them against each other. Equation (2) divides the shock between quantity and price according to the elasticity of local labour supply, and the division observed here, with employment clearly up at 9.237 and the aggregate hourly wage up but imprecise at 4.985 (s.e. 5.951), matches what an elastic local supply produces. Exposure to AI-adopting industries was, over this decade, a positive local labour-demand shock whose incidence favoured men. The next paragraphs stress-test that fact; the rest of the section then asks which channel carries it to the marriage market.

Table 6: AI exposure and earnings by sex, 2010–2021

	Δ Mean ln earnings, men			Δ Mean ln earnings, women			Δ Earnings gap (M−F)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	-0.124	7.068	3.449	-4.611	2.955	-1.939	4.487*	4.113	5.388
	(4.780)	(4.693)	(3.836)	(3.738)	(3.445)	(2.730)	(2.607)	(3.062)	(3.692)
Female share		-2.351***	-1.852***		-2.179***	-1.584***		-0.172	-0.268
		(0.539)	(0.484)		(0.506)	(0.414)		(0.297)	(0.359)
Age 60+ share		0.715*	0.538		0.921**	0.856**		-0.206	-0.318
		(0.423)	(0.430)		(0.349)	(0.324)		(0.382)	(0.369)
BA+ share			-0.310**			-0.186			-0.124
			(0.131)			(0.111)			(0.0830)
High-AI occ., women			1.081***			0.399			0.682***
			(0.323)			(0.286)			(0.251)
High-AI occ., men			0.0742			0.478**			-0.404**
			(0.211)			(0.197)			(0.169)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.000	0.081	0.136	0.009	0.098	0.157	0.009	0.011	0.032
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	5.969	15.51***	15.42***	-3.571	4.724	-4.240	9.540***	10.79***	19.66***
	(5.561)	(5.417)	(5.395)	(3.930)	(4.048)	(4.040)	(3.000)	(3.548)	(5.373)
Female share		-2.734***	-2.058***		-2.259***	-1.544***		-0.475	-0.514
		(0.535)	(0.463)		(0.534)	(0.418)		(0.306)	(0.387)
Age 60+ share		1.112***	0.812*		1.004***	0.803**		0.108	0.00901
		(0.431)	(0.443)		(0.367)	(0.320)		(0.371)	(0.351)
BA+ share			-0.439***			-0.161			-0.279***
			(0.134)			(0.112)			(0.0951)
High-AI occ., women			1.011***			0.412			0.599**
			(0.297)			(0.282)			(0.242)
High-AI occ., men			0.0727			0.478**			-0.406**
			(0.201)			(0.195)			(0.161)
Observations	741	741	741	741	741	741	741	741	741
R-squared	-0.013	0.062	0.113	0.008	0.097	0.156	-0.003	-0.004	-0.011

CZ-level long differences in the mean of log annual wage and salary income among 18–64 year olds with positive earnings. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$. Columns (1), (4) and (7) include no controls; (2), (5) and (8) add the baseline female and age 60+ shares; (3), (6) and (9) add the baseline BA+ share and the sex-specific baseline shares of employment in AI-exposed occupations.

Timing and confounders. The labour-market evidence passes the same checks as the marital estimates, and the two confounders bear on it directly. Table 7 runs every labour-market outcome over the four windows of Table 3. Over 2000–2007, before diffusion began, exposure predicts nothing: employment is 2.972 (s.e. 2.956) and the earnings gap 2.106 (s.e. 5.593). The one exception is the dispersion of annual earnings, which was already falling differentially in exposed places before 2007; dispersion plays no role in the mechanism the results support. The 2007–2010 column quantifies the recession: employment (-5.624), male full-time work (-11.132) and male earnings (-18.064) all fell differentially and precisely, so part of the post-2010 expansion is recovery.

Table 7: AI exposure and labour-market outcomes across four windows: pre-recession placebo, recession, pre-COVID and full sample

	Placebo	**Recession**	**Treatment window**	
	2000–2007	2007–2010	2010–2019	2010–2021
Δ Employment rate	2.972	-5.624**	9.508***	9.237***
	(2.956)	(2.221)	(2.586)	(2.309)
Δ Full-time rate, men	3.623	-11.13***	13.26***	13.50***
	(3.978)	(3.405)	(3.383)	(4.484)
Δ Full-time gap (M−F)	2.429	-5.137***	7.787***	8.055***
	(2.404)	(1.868)	(2.231)	(2.748)
Δ Mean ln earnings, men	6.939	-18.06***	7.270	15.42***
	(11.07)	(6.214)	(4.669)	(5.395)
Δ Mean ln earnings, women	4.832	-14.21***	-4.986	-4.240
	(6.118)	(3.941)	(5.813)	(4.040)
Δ Earnings gap (M−F)	2.106	-3.859	12.26***	19.66***
	(5.593)	(4.739)	(4.570)	(5.373)
Δ S.d. of ln earnings	-10.04**	-1.005	-1.550	4.224
	(4.819)	(2.555)	(4.017)	(5.145)
Δ Zero-earnings share	-1.830	6.083***	-8.167***	-8.270***
	(2.634)	(2.178)	(2.538)	(2.011)
Δ AI-exposed employment share	-0.580	-0.942	2.743**	6.010***
	(0.969)	(1.017)	(1.126)	(1.218)

Each cell is a separate 2SLS long-difference regression of the row outcome on AI exposure, with the full control set measured at the base year of that window, so the 2000–2007 column conditions on nothing realised after 2000. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Restricting the endline to 2019 leaves the results intact (employment is 9.508 and the earnings gap 12.256), and Figure 10 makes the timing visible directly. Holding the 2010 base fixed and moving the endpoint year from 2011 through 2021, the employment, full-time-gap and earnings-gap coefficients rise steadily from 2012 and are already large and statistically significant by 2016. They do not jump in 2020, so the estimates do not originate with the pandemic. Nor is the response a tail phenomenon: when the linear term is replaced with population-weighted quintiles of exposure, the employment rate and the earnings gap rise across every quintile and closely track the linear fit (Figure 11). This dose–response pattern indicates that the labour-demand channel operates across the range of exposure, and not only in a handful of unusual commuting zones.

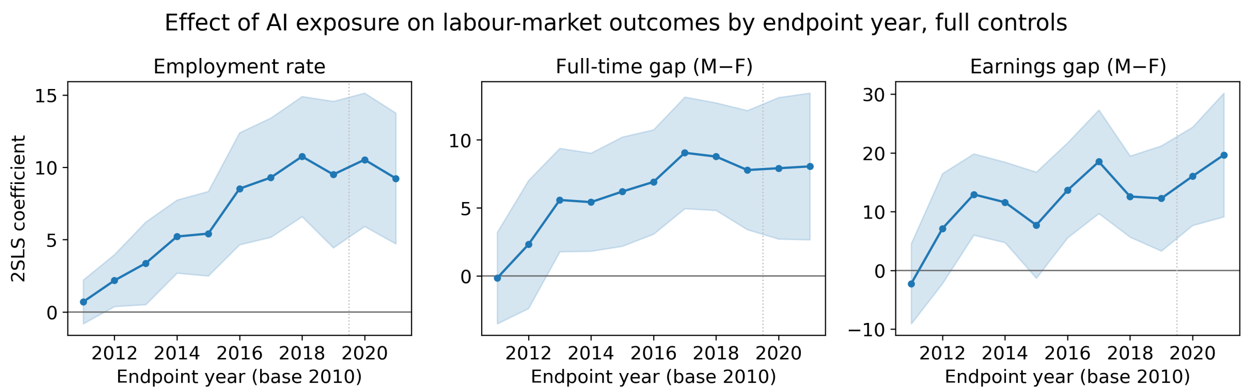

Figure 10: **Effect of AI exposure on labour-market outcomes by endpoint year.** Each point is the fully controlled 2SLS coefficient from a separate regression of the 2010-to-t long difference on AI exposure; shaded bands are 95 percent confidence intervals from state-clustered standard errors. The dotted vertical line marks the last pre-pandemic endpoint.

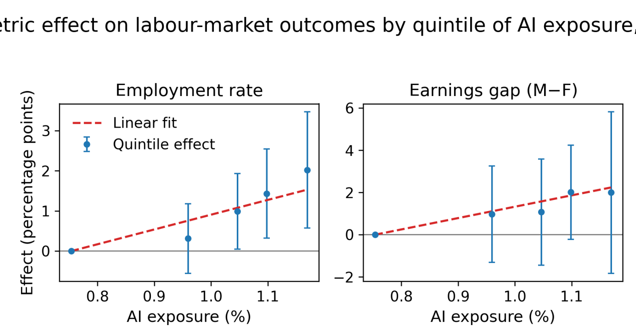

Figure 11: **Non-parametric effect of AI exposure on labour-market outcomes.** Points are the estimated effects of population-weighted quintiles of AI exposure relative to the lowest quintile, from OLS regressions with the full control set, plotted at the population-weighted mean exposure of each quintile; bars are 95 percent confidence intervals from state-clustered standard errors. The dashed line is the linear fit from the corresponding single-coefficient regression. Effects are in percentage points.

The recession is the confounder this timing evidence cannot dispose of, since recovery also builds through the 2010s. Table 8 conditions on each commuting zone's own 2007–2010 changes, as Table 4 does for the marital outcomes. Employment falls from 9.237 to 3.970 and the earnings gap from 19.662 to 11.778, both still significant at the 5 percent level. Recovery therefore explains part of the labour-market response without eliminating it.

Table 8: Labour-market outcomes controlling for 2007–2010 recession severity, 2010–2021

	Δ Employment rate			Δ Earnings gap (M−F)		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS Estimates						
AI Exposure	3.681**	1.959**	1.815**	5.388	4.041	0.957
	(1.565)	(0.888)	(0.832)	(3.692)	(3.440)	(3.173)
Δ_{07}^{10} Employment rate		-0.728***	-0.701***		-0.569**	-0.307
		(0.0589)	(0.0867)		(0.276)	(0.346)
Δ_{07}^{10} Full-time rate, men			-0.0108			0.105
			(0.0563)			(0.261)
Δ_{07}^{10} Mean ln earnings, men			-0.0126			-0.350***
			(0.0189)			(0.0780)
Panel B: 2SLS Estimates						
AI Exposure^{2sls}	9.237***	4.125**	3.970**	19.66***	16.41***	11.78**
	(2.309)	(1.773)	(1.770)	(5.373)	(5.290)	(4.851)
Δ_{07}^{10} Employment rate		-0.707***	-0.693***		-0.449*	-0.267
		(0.0601)	(0.0831)		(0.268)	(0.349)
Δ_{07}^{10} Full-time rate, men			-0.00468			0.135
			(0.0550)			(0.256)
Δ_{07}^{10} Mean ln earnings, men			-0.00767			-0.325***
			(0.0176)			(0.0745)

Counterpart of Table 4 for the labour-market outcomes: columns add the commuting zone's own 2007–2010 changes in employment, male full-time work and male earnings as controls, so the comparison is between places that experienced the recession similarly. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Earnings risk and displacement. Table 9 tests the first two requirements of the risk channel. The within-commuting-zone standard deviation of log earnings does not move: the fully controlled 2SLS estimate is 4.224 with a standard error of 5.145. The share of adults with zero earnings falls by -8.270, and the share of employment in occupations above the 66^{th} percentile of the Webb distribution, precisely the workers the index identifies as most exposed to replacement, rises by 6.010; both are significant at the 1 percent level. Over 2010–2021 and at this level of aggregation, the more AI-adopting commuting zones employed relatively more of their AI-exposed workers.

The dispersion result does not survive a better measure. The standard deviation in Table 9 is computed over annual earnings, which hours movements contaminate. Recomputed over hourly wages it rises by 6.793, significant at the 1 percent level, by 5.196 among full-time full-year workers, and by more for men (7.179) than for women (4.966). The 2000–2007 placebo is -4.406 and indistinguishable from zero, so the rise is specific to the window in which AI diffused. On the measure that holds hours fixed, then, the first requirement of the risk channel is met: local wage dispersion did widen where exposure was higher.

The widening changes the basis on which the risk channel is rejected. A cross-sectional standard deviation across the workers of a commuting zone is not the object the model calls σ_H^2, which is the variance of the shock facing one worker. The two coincide only when the widening reflects greater unpredictability; a greater spread of predictable returns also widens the cross-section, and a demand expansion that raises pay at the top of the local distribution does exactly that while leaving any individual's earnings no harder to forecast. That reading fits the rest of the evidence: dispersion widens most for men, whose mean wage rises, and least for women, whose mean wage does not move. The design cannot separate the two interpretations, so the dispersion evidence is consistent with the risk channel and does not discriminate against it. The discriminating evidence lies in the other two requirements: exposure moves employment into the occupations the risk channel says it should move workers out of, and the couple-level evidence below shows none of the tilt toward mixed couples that channel requires.

Table 9: AI exposure, earnings dispersion and displacement, 2010–2021

	Δ S.d. of ln earnings			Δ Zero-earnings share			Δ AI-exposed employment share		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	10.49***	6.388**	2.119	-2.326**	-2.789*	-3.190**	5.896***	5.801***	2.474***
	(3.030)	(2.922)	(2.828)	(1.061)	(1.433)	(1.505)	(0.514)	(0.803)	(0.786)
Female share		0.858	0.0957		-0.161	-0.0932		0.116	-0.131
		(0.621)	(0.444)		(0.196)	(0.208)		(0.0982)	(0.0981)
Age 60+ share		-0.685**	-0.365*		-0.225**	-0.204*		0.0395	0.123
		(0.261)	(0.198)		(0.109)	(0.103)		(0.115)	(0.0980)
BA+ share			0.570***			0.0141			0.166***
			(0.104)			(0.0280)			(0.0328)
High-AI occ., women			-1.079***			-0.154*			-0.0390
			(0.229)			(0.0852)			(0.0653)
High-AI occ., men			-0.0665			0.133**			-0.0953**
			(0.145)			(0.0604)			(0.0445)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.073	0.106	0.231	0.031	0.057	0.078	0.209	0.213	0.323
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	11.42***	6.997	4.224	-3.589**	-5.107**	-8.270***	7.366***	8.039***	6.010***
	(2.870)	(4.914)	(5.145)	(1.525)	(1.998)	(2.011)	(0.661)	(0.890)	(1.218)
Female share		0.830	0.0594		-0.0558	-0.00577		0.0151	-0.191*
		(0.733)	(0.466)		(0.180)	(0.188)		(0.112)	(0.104)
Age 60+ share		-0.656*	-0.317		-0.334***	-0.320***		0.145	0.204**
		(0.376)	(0.254)		(0.110)	(0.114)		(0.116)	(0.0966)
BA+ share			0.547***			0.0692**			0.128***
			(0.118)			(0.0304)			(0.0392)
High-AI occ., women			-1.092***			-0.125			-0.0595
			(0.225)			(0.0814)			(0.0626)
High-AI occ., men			-0.0668			0.134**			-0.0957**
			(0.143)			(0.0574)			(0.0422)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.072	0.106	0.230	0.022	0.034	0.010	0.196	0.190	0.289

CZ-level long differences, 2010–2021. The standard deviation of log earnings is computed across 18–64 earners within the commuting zone; the zero-earnings share is measured among those in the wage universe; the AI-exposed employment share is the share of employment in occupations above the 66^{th} percentile of the Webb (2019) index. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$. Columns (1), (4) and (7) include no controls; (2), (5) and (8) add the baseline female and age 60+ shares; (3), (6) and (9) add the baseline BA+ share and the sex-specific baseline shares of employment in AI-exposed occupations.

Discussions of AI and the labour market often assume a displacement premise without demonstrating it; over this window it fails. Whether the same holds of the generative-AI wave that began after this sample ends is a question these data cannot answer: the ABS measures machine learning, machine vision, natural language processing, voice recognition and automated guided vehicles as deployed by 2020, a different technology from the one now diffusing.

Incidence by sex. Exposure does move the position of men relative to women. Table 10 reports each margin separately for men, for women and for the gap between them. Exposure does not raise men's position by pushing women out of work: the employment rate rises for both sexes, by 9.974 for men and 8.415 for women, and the gap between the two is indistinguishable from zero. The sexes diverge on hours and on earnings. The full-time rate rises by 13.503 for men against 5.448 for women, widening the full-time gap by 8.055, and Table 6 showed the male–female earnings gap widening by 19.662. Under the classical specialisation logic (Becker, 1973; Autor et al., 2019), a shock that lowers men's relative standing in the labour market lowers their value in the marriage market and reduces marriage, and a shock that raises their standing should therefore raise it.

Table 10: AI exposure by sex: employment, hours and hourly wages, 2010–2021

	Men	Women	Gap (M−F)
Employment rate	9.974***	8.415***	1.558
	(3.353)	(1.994)	(2.980)
Full-time rate	13.50***	5.448**	8.055***
	(4.484)	(2.561)	(2.748)
Mean ln annual earnings	15.42***	-4.240	19.66***
	(5.395)	(4.040)	(5.373)
Hourly wages (weeks-worked extract)			
Hourly wage	10.73	-1.513	12.24***
	(6.885)	(5.593)	(3.692)
Hourly wage, full-time full-year	10.02	-1.408	11.42**
	(9.598)	(7.510)	(4.538)
Hourly wage, flat 52-week denominator	10.74**	-3.741	14.48***
	(5.456)	(5.143)	(4.505)
The hourly wage under the paper's two threats			
2000–2007 placebo	3.157	3.749	-0.592
	(6.298)	(3.853)	(3.258)
2007–2010 recession controls	6.772	-3.003	9.775**
	(6.416)	(4.887)	(3.914)

Each cell is a separate 2SLS long-difference regression of the row outcome on AI exposure with the full control set, estimated for men, for women and for the male–female gap. The employment, full-time and annual-earnings rows use the main ACS panels of Tables 5 and 6. The hourly rows divide annual wage and salary income by usual weekly hours times weeks worked in the year; the ACS reports weeks worked in intervals, so annual hours use the interval midpoints, and records outside the weeks universe are dropped rather than imputed. The flat 52-week row divides by fifty-two weeks on the same sample. The placebo row re-estimates the hourly outcomes over 2000–2007 with controls measured in 2000; the recession row adds the commuting zone's own 2007–2010 changes in employment, male full-time work and male earnings to the controls. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

The two margins are not equally robust. Both are absent in the 2000–2007 placebo, at 2.106 for the earnings gap and 2.429 for the full-time gap. Conditioning on each commuting zone's own 2007–2010 experience, the earnings gap survives at 11.778 while the full-time gap falls to 2.930 and is no longer distinguishable from zero. The hours margin is therefore substantially recovery from the recession trough in places that had fallen furthest, and I do not lean on it. The incidence claim rests on the earnings gap.

The gap is not merely an hours effect. The outcome in Table 6 is annual wage and salary income, which is hours multiplied by the wage, and men's full-time rate rises by two and a half times as much as women's, so some of the widening has to be hours. To separate the two I build a wage per hour, dividing annual earnings by usual weekly hours times weeks worked in the year; the hourly rows of Table 10 report the estimates. On that measure the gap is 12.241 against 19.662 for annual earnings, significant at the 1 percent level, so close to two thirds of it survives with hours held fixed. Men's hourly wage rises by 10.728, women's is flat at -1.513, and among workers employed full time for the full year the gap is 11.423. It behaves under the paper's two main threats as the annual gap does: it is -0.592 in the 2000–2007 placebo, and survives at 9.775 conditioning on each commuting zone's own 2007–2010 experience. Exposure moved men's pay per hour, and not only the hours they worked.

The ACS reports weeks worked in intervals, so annual hours use the interval midpoints, and records outside the weeks universe are dropped without imputation. Dividing instead by a flat fifty-two weeks puts the gap at 14.477 on the same sample. The overstatement is entirely on the women's side: a flat denominator understates the hourly wage of part-year workers, exposure draws women into employment, and the entrants it draws in are disproportionately part-year. Measuring their weeks removes the overstatement, which is the reason to prefer the smaller number.

The incidence does not come from the sex composition of the adopting industries. That composition runs the other way. The industries with the highest adoption rates include machinery, fabricated metal products, computer and electronic products and transportation equipment, which score highly on machine vision and automated guided vehicles and are 77 percent male. They employ 5.1 million people between them and carry 8.3 percent of the shifter. Health care and education employ 29.4 million, are 26 percent male, and carry 28.6 percent of it. When industries are weighted by employment instead of each adoption rate counting alike, the correlation between adoption and male intensity is -0.173, and the average woman works in industries adopting at 1.162 percent against 1.067 percent for the average man. Whatever produces the earnings gap therefore operates within industries or within occupations, and a treatment that varies only across industries cannot decompose it. The design signs and sizes the incidence without explaining it.

This is the same mechanism documented by Autor et al. (2019) and Anelli et al. (2024), operating in the opposite direction. Trade and robots lowered men's relative earnings and employment in exposed places, and marriage fell; AI exposure raised them, and marriage rose. The difference from earlier automation therefore lies in the sign and incidence of the labour-demand shock that accompanied diffusion over this particular decade, and not in the nature of the technology.

Income effects versus insurance. A third possibility is observationally close to the insurance channel: agents may marry to share the fixed costs of a household when they become poorer, so that marriage rises where earnings fall. This predicts the opposite sign of what is observed. Exposure raises male earnings and the employment rate, so a poverty-driven cost-sharing account points the wrong way for the marriage response documented in Table 2; the level of earnings rejects it on its own, with no appeal to dispersion. The remaining choice is between the two channels of Section 3, and the couple-level evidence below completes it.

Couple-level exposure and assortative matching. Table 11 decomposes the flow of first marriages by the occupational exposure profile of the couple, with both partners' occupations frozen at 2010. Entries are categorised as low–low (both partners at or below the 33rd percentile of the Webb distribution), mixed (one at or below the 33rd and one at or above the 66th), and high–high (both at or above the 66th); Appendix Table B13 reports summary statistics for the three shares. The low–low estimate, for couples the shock does not reach, is 0.322 (s.e. 0.382). Mixed and high–high couples gain strongly, 2.127 and 2.055 respectively, and the high–high gain exceeds the low–low one by 1.734, significant at the 1 percent level. Relative to each type's own 2010 population share, the gains rise with the number of exposed partners: about a tenth of the base share for low–low, four tenths for mixed and two thirds for high–high.

Under a risk-sharing reading the high–high gains would be puzzling, since pairing two positively correlated risks is a poor way to insure against either, and part (iii) of Proposition 1 makes the point exact. The proposition predicts the sign of a cross-difference; the ordering of the three coefficients is not restricted. The earnings term is modular, so under the labour-demand channel $S_{HH} + S_{LL} - S_{HL} - S_{LH}$ does not move; the risk term is strictly submodular for any positive spousal correlation, so a shock working through risk has to drive that combination down, placing its marriages in mixed couples at the expense of both same-type couples. Estimating the cross-difference directly gives 0.250 (s.e. 0.569), zero to sampling precision, exactly as modularity requires, and the log-odds counterpart from Equation (12) is equally indistinguishable from zero. That comparison holds for every value of θ and every $\rho > 0$, so it bears on the channel even though the sign of the risk term itself is not pinned down, and it shows none of the tilt toward mixed couples the risk channel needs. Part (ii) of the same Proposition supplies the reading the data support: sorting on earnings capacity, in which a labour-demand expansion concentrated in AI-adopting industries raises the earnings of workers in AI-exposed occupations and marriage among them rises accordingly. One caveat on measurement applies. This exercise interacts a commuting-zone-level treatment with an individual-level occupational index, so that AI exposure appears on both sides of the equation. These patterns should therefore be read as descriptive evidence consistent with the mechanism. They are not independent causal estimates.

Table 11: AI exposure and first marriages by couple type, 2010–2021

	Δ Share low–low 1^{st} marriages			Δ Share mixed 1^{st} marriages			Δ Share high–high 1^{st} marriages		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	0.0503	-0.0135	0.0339	0.688***	0.542*	0.710**	2.124***	2.214***	1.004***
	(0.168)	(0.215)	(0.286)	(0.254)	(0.287)	(0.307)	(0.275)	(0.418)	(0.313)
Female share		-0.00690	-0.0350		0.0498*	-0.0161		-0.0248	-0.0344
		(0.0286)	(0.0295)		(0.0293)	(0.0301)		(0.0338)	(0.0357)
Age 60+ share		-0.0223	-0.0178		-0.0134	-0.00970		0.0115	0.0343
		(0.0227)	(0.0242)		(0.0273)	(0.0295)		(0.0499)	(0.0451)
BA+ share			0.0111			0.0174			0.0347**
			(0.00937)			(0.0113)			(0.0157)
High-AI occ., women			-0.0108			0.0131			-0.0296
			(0.0222)			(0.0251)			(0.0329)
High-AI occ., men			-0.0206			-0.0662***			0.0396**
			(0.0141)			(0.0162)			(0.0167)
Observations	741	741	741	741	741	741	741	741	741
R-squared	0.000	0.003	0.018	0.021	0.025	0.074	0.180	0.181	0.262
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	0.128	0.0514	0.322	1.045***	1.056***	2.127***	2.751***	3.110***	2.055***
	(0.183)	(0.256)	(0.382)	(0.318)	(0.377)	(0.452)	(0.267)	(0.415)	(0.452)
Female share		-0.00984	-0.0400		0.0266	-0.0405		-0.0654*	-0.0525*
		(0.0292)	(0.0270)		(0.0276)	(0.0291)		(0.0387)	(0.0314)
Age 60+ share		-0.0192	-0.0112		0.0109	0.0227		0.0537	0.0584
		(0.0244)	(0.0236)		(0.0276)	(0.0282)		(0.0495)	(0.0437)
BA+ share			0.00797			0.00209			0.0233
			(0.0106)			(0.0127)			(0.0191)
High-AI occ., women			-0.0124			0.00488			-0.0357
			(0.0215)			(0.0251)			(0.0323)
High-AI occ., men			-0.0206			-0.0663***			0.0395**
			(0.0138)			(0.0164)			(0.0164)
Observations	741	741	741	741	741	741	741	741	741
R-squared	-0.000	0.003	0.015	0.015	0.016	0.034	0.165	0.158	0.242

CZ-level long differences in the share of 18–64 year olds in a first marriage of each couple type, classified by both partners' 2010 occupations against the Webb (2019) AI-exposure distribution (low $\leq$ 33rd percentile, high $\geq$ 66th). All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$. Columns (1), (4) and (7) include no controls; (2), (5) and (8) add the baseline female and age 60+ shares; (3), (6) and (9) add the baseline BA+ share and the sex-specific baseline shares of employment in AI-exposed occupations.

Do the occupationally exposed work in the adopting industries? The commuting-zone design carries a second intention-to-treat assumption alongside the standard shift-share one: that the workers an occupational exposure index flags are in fact the workers whose industries go on to adopt AI. If the two were unrelated, the industry-level treatment would have no claim on the population the occupational index identifies. Table 12 puts both variables on the same person, using the 1,280,829 employed 18–64 year olds in the 2010 ACS for whom an occupational Webb score and a mapped industry are both available.

The link holds. Mean industry adoption rises monotonically across quartiles of the worker's own occupational exposure, from 0.978 percent in the least exposed quartile to 1.387 percent in the most exposed, with a worker-level correlation of 0.212. Workers above the 66^{th} percentile of the Webb distribution were employed in industries adopting at 1.361 percent against 1.024 percent for everyone else, a difference of 0.310 percentage points once sex, education and age are held constant. The alignment is far from perfect, as one would expect of an occupational index and an industry survey, but it runs in the right direction and is monotone.

The alignment establishes that the industry treatment does reach the workers the occupational index identifies as exposed. The next result, that those workers show no differential response through their own industry, is therefore informative rather than mechanical.

Table 12: Do workers in AI-exposed occupations work in the industries that adopt AI?

	Share of employment (%)	Industry AIAdopt (%)	Industry ML adoption (%)
Q1 (lowest)	26.7	0.978	1.443
Q2	24.9	1.042	1.539
Q3	25.9	1.101	1.798
Q4 (highest)	22.5	1.387	2.113
Correlation with own Webb percentile			
across all workers		+0.212	+0.197

Employed 18–64 year olds in the 2010 ACS with both an occupational Webb (2019) score and a mapped industry (1,280,829 person-records, person-weighted). Rows are quartiles of the worker's own occupational AI-exposure score; columns report the AI adoption rate of the industry that worker was employed in, taken from the ABS Technology Module. Industry is assigned from the census industry classification through a crosswalk to the 46 ABS industries, which places 93.1 percent of employment. Workers above the 66^{th} percentile of the Webb distribution were in industries adopting at 1.361 percent against 1.024 percent for the rest; conditioning on sex, education and age the difference is 0.310 percentage points (s.e. 0.002).

Within-commuting-zone evidence. A final question is whether the effect operates through the workers whose own industries adopted AI, or through the local labour market as a whole. Equation (17) implies the latter, and the distinction is testable. Table 13 compares industries within the same commuting zone: cells are commuting zone $\times$ industry, the regressor is that industry's ABS adoption rate, and commuting-zone fixed effects absorb every local shock (housing, migration, policy, recession severity, pandemic severity).

There is no differential response. Marriage is -0.228 (s.e. 0.482), first marriage 0.106 (s.e. 0.429), the divorced-to-married ratio -0.075 (s.e. 0.453), earnings 0.105 (s.e. 2.047) and earnings dispersion -1.591 (s.e. 1.555), across 10,689 cells in 741 commuting zones and 43 industries. The comparison is informative precisely because the two coefficients are on the same scale: the commuting-zone regressor is the share-weighted average of exactly these industry adoption rates, so if the aggregate estimate were the aggregation of own-industry effects the two would coincide. The first-marriage estimate of 5.758 lies far outside the confidence interval around 0.106, so an own-industry mechanism of the magnitude required to generate the aggregate result is rejected.

Table 13: Within commuting zone: workers in AI-adopting industries

	2010–2021	2010–2019
Δ Share married	-0.228	0.00949
	(0.482)	(0.321)
Δ Share 1^{st} married	0.106	0.351
	(0.429)	(0.309)
Δ Share divorced / married	-0.0747	-0.556
	(0.453)	(0.416)
Δ Mean ln earnings	0.105	-1.348
	(2.047)	(1.440)
Δ S.d. of ln earnings	-1.591	-1.044
	(1.555)	(1.190)
Cells	10689	10825
Commuting zones	741	741
Industries (clusters)	43	43

Each cell is a separate regression of the long difference in the row outcome, measured over the workers of one industry in one commuting zone, on that industry's ABS AI adoption rate, with commuting-zone fixed effects absorbing every local shock. The sample is 18–64 year olds who worked in the previous year, with industry taken from the census industry classification; cells need at least 30 person-records at both endpoints. Regressions are weighted by cell population and standard errors are clustered by industry, the level at which the shocks vary. Because the shift-share regressor is the share-weighted average of the same industry adoption rates, these coefficients are directly comparable with the commuting-zone estimates in Table 2.

A labour-demand interpretation implies this pattern, and a worker-level displacement account does not. A demand expansion in some local industries raises employment, earnings and marriage across the local labour market, including for workers whose own industries adopted nothing. The design also addresses the ecological-inference concern that a commuting-zone treatment cannot support conclusions about individual behaviour: the individual-level version of the claim is tested directly here and does not hold, so the aggregate claim is restricted to a statement about local labour markets. Two caveats bound the reading. These cells condition on employment, and employment itself responds to exposure, so cell composition shifts; and attenuation from sampling error in small cells biases the coefficients toward zero, so the null is better read as an upper bound on the own-industry channel than as a precise zero.

6 Conclusion

This paper has investigated how the diffusion of artificial intelligence relates to marriage and divorce across U.S. local labour markets. Combining a Bartik-style measure of commuting-zone AI exposure, built from the ABS Technology Module and instrumented with its European analogue, with ACS microdata on marital and labour market outcomes, it has documented that commuting zones more specialised in industries that went on to adopt AI experienced relatively more first marriages and relatively fewer divorces over 2010–2021.

The theoretical framework shows that a sectoral shock can raise the marital surplus either by raising mean earnings in the exposed sector or by raising their variance, and that the two carry different empirical signatures. A shock working through the level of earnings raises employment and men's relative earnings, need not touch the dispersion of earnings, and places the marriages it creates in proportion to how exposed both partners are. A shock working through risk has to raise dispersion, has to push workers out of the exposed occupations, and has to place its marriages in couples whose partners are differently exposed, the last of these for every value of the household parameters.

The data select decisively between them. AI exposure raises the employment rate, lowers the share of adults with no earnings, raises the share of employment in the very occupations most exposed to AI, and adds marriages in proportion to the couple's exposure. Local wage dispersion does widen, but a demand expansion that raises pay at the top of the distribution widens it just as an increase in earnings risk would, so the rejection of the risk channel rests on displacement and on where the marriages form. Exposure does move the earnings and full-time employment of men relative to women. Because the household converts the higher earner's wage into consumption at the higher rate, the framework predicts marriage to rise under exactly this incidence, and the marriage response follows the sign that the literature on manufacturing decline would predict from a shock of that incidence. The incidence does not, however, derive from which industries adopted: measured by employment, adoption over this window is concentrated slightly more in the industries where women work, so whatever raises men's relative position operates within industries, and the design signs it without explaining it.

That literature has consistently found that technology and trade shocks reduce marriage. The results here show that the sign follows from the accompanying labour-demand shock and from how that shock falls across the sexes. Robots and import competition fell on male-intensive industries and lowered male relative earnings; AI adoption over this decade arrived in industries in which employment rose and men gained relative ground, and the demographic response followed.

Four limitations bound the claim. First, the design identifies the effect of specialising in AI-adopting industries, and carries no claim about the effect of AI holding constant everything correlated with adoption: with 13.5 effective shocks, a random reshuffling of the adoption vector performs comparably to the true one. Second, the estimates are differential across commuting zones and are silent about the national average effect. Third, the sex incidence on which the mechanism turns is measured but not decomposed, since it does not come from the industry mix and the treatment varies only across industries; the part of it that is a movement into full-time work is not separable from recovery after 2010, though the hourly-wage component is. Fourth, the sample ends in 2021, before the generative-AI wave; whether the labour-demand incidence documented here survives that transition is an open question, and if the sign of the shock changes, the framework implies that the demographic response should change with it.

On the policy side, the results suggest that the demographic consequences of a technology wave are mediated by which workers it complements. Policies that shape the incidence of technological change across sectors and sexes will therefore have consequences for family formation that are separate from, and potentially larger than, their direct labour market effects.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from us commuting zones. *American Economic Review: Insights*, 4(4):435–54.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from u.s. labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Adão, R. R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Economic policy for artificial intelligence. *Innovation Policy and the Economy*, 19:139–159.

Aguiar, M. and Hurst, E. (2007). Measuring trends in leisure: The allocation of time over five decades. *The Quarterly Journal of Economics*, 122(3):969–1006.

Anelli, M., Giuntella, O., and Stella, L. (2024). Robots, marriageable men, family, and fertility. *Journal of Human Resources*.

Autor, D. H., Dorn, D., and Hanson, G. H. (2019). When work disappears: Manufacturing decline and the falling marriage-market value of young men. *American Economic Review: Insights*, 1(2):161–178.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.

Becker, G. S. (1973). A theory of marriage: Part i. *Journal of Political Economy*, 81(4):813–846.

Becker, G. S. (1981). *A Treatise on the Family*. Harvard University Press, Cambridge, MA.

Black, S. E. and Spitz-Oener, A. (2010). Explaining women's success: Technological change and the skill content of women's work. *Review of Economics and Statistics*, 92(1):187–194.

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *American Economic Review*, 106(2):387–435.

Borusyak, K., Hull, P., and Jaravel, X. (2023). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 90(2):883–928.

Browning, M., Chiappori, P.-A., and Lewbel, A. (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies*, 80(4):1267–1303.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129(4):1553–1623.

Chiappori, P.-A., Costa Dias, M., and Meghir, C. (2018). The marriage market, labor supply, and education choice. *Journal of Political Economy*, 126(S1):S26–S72.

Chiappori, P.-A., Salanié, B., and Weiss, Y. (2017). Partner choice, investment in children, and the marital college premium. *The American Economic Review*, 107(8):2109–2167.

Choo, E. and Siow, A. (2006). Who marries whom and why. *Journal of Political Economy*, 114(1):175–201.

Costanzo, C. (2025). Robots, jobs, and optimal fertility timing. *Journal of Population Economics*, 38(2):1–33.

Dingel, J. I. and Neiman, B. (2020). How many jobs can be done at home? *Journal of Public Economics*, 189:104235.

Eika, L., Mogstad, M., and Zafar, B. (2019). Educational assortative mating and household income inequality. *Journal of Political Economy*, 127(6):2795–2835.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

Galichon, A. and Salanié, B. (2022). Cupid's invisible hand: Social surplus and identification in matching models. *The Review of Economic Studies*, 89(5):2600–2629.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Greenwood, J. and Guner, N. (2004). Marriage and divorce since world war ii: Analyzing the role of technological progress on the formation of households. NBER Working Paper 10772, National Bureau of Economic Research.

Greenwood, J., Seshadri, A., and Yorukoglu, M. (2005). Engines of liberation. *Review of Economic Studies*, 72(1):109–133.

Hess, G. D. (2004). Marriage and consumption insurance: What's love got to do with it? *Journal of Political Economy*, 112(2):290–318.

Hyslop, D. R. (2001). Rising US earnings inequality and family labor supply: The covariance structure of intrafamily earnings. *The American Economic Review*, 91(4):755–777.

Kotlikoff, L. J. and Spivak, A. (1981). The family as an incomplete annuities market. *Journal of Political Economy*, 89(2):372–391.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25 (NeurIPS)*.

Matysiak, A., Bellani, D., and Bogusz, H. (2023). Industrial robots and regional fertility in european countries. *European Journal of Population*, 39(11):1–28.

Moretti, E. (2011). Local labor markets. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 4B, pages 1237–1313. Elsevier.

Ngai, L. R. and Petrongolo, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4):1–44.

Rosenzweig, M. R. and Stark, O. (1989). Consumption smoothing, migration, and marriage: Evidence from rural india. *Journal of Political Economy*, 97(4):905–926.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Rivera Drew, J. A., Richards, S., Rogers, R., Schroeder, J., and Williams, K. C. W. (2025). Ipums usa: Version 16.0. Minneapolis, MN. Dataset.

Schaller, J. (2016). Booms, busts, and fertility: Testing the becker hypothesis. *Journal of Human Resources*, 51(1):1–46.

Webb, M. (2019). The impact of artificial intelligence on the labor market.

Wilson, W. J. (1996). *When Work Disappears: The World of the New Urban Poor*. Knopf, New York.

Acknowledgements

Data availability A replication package containing all code and derived datasets reproduces every table and figure in this paper from public sources. The IPUMS extract is not redistributed under the IPUMS terms of use; the package documents how to recreate it.

Declarations

Conflict of interest The author declares that he has no conflict of interest.

Appendix

A Proofs

Proof of Proposition 1. For (i), $t_s(w) = (\kappa\alpha/w)^\nu$ is strictly decreasing in w, and $2\theta^{\nu-1}-1 > 1$ because $\theta > 1$ and $\nu > 1$. The right-hand expression in Equation (10) is therefore strictly increasing in w; evaluating it at w^m and at w^f and differencing gives the displayed wedge, which is positive because $w^m > w^f$. Rearranging $D'(w) > 0$ gives the stated condition, which has a solution only if $\theta < 2$, since the right-hand side of Equation (10) is bounded above by $2/\theta - 1$.

For (ii), S is the sum of a function of w^m alone, a function of w^f alone and a term in the variances, so the earnings component has zero cross-partial in the two wages and the displayed difference cancels term by term. Let the exposed class's wages rise by Δw_H and the unexposed class's by $\Delta w_L < \Delta w_H$. Then $S(H, H)$ gains $[D'(w^m_H) + D'(w^f_H)]\Delta w_H$ and $S(L, L)$ gains $[D'(w^m_L) + D'(w^f_L)]\Delta w_L$, with each mixed match taking one term from each; since $D' > 0$ under the condition in (i) and D' is increasing, the ordering follows.

For (iii), write R for the risk term in Equation (8). Because shocks are correlated at ρ within a class and orthogonal across classes, $R(H,H) = \gamma\sigma_H^2[1-2(1+\rho)/\theta^2]$, $R(L,L) = \gamma\sigma_L^2[1-2(1+\rho)/\theta^2]$ and $R(H,L) = R(L,H) = (\gamma/2)(\sigma_H^2+\sigma_L^2)-(\gamma/\theta^2)(\sigma_H^2+\sigma_L^2)$. Summing, the cross-difference is $-(2\gamma\rho/\theta^2)(\sigma_H^2+\sigma_L^2)$, and differentiating in σ_H^2 gives the result. $\square$

B Supplementary tables

Table B1: Descriptive statistics: outcome changes (2010–2021)

Variable ($\Delta_{2010-2021}$)	N	Mean	Std. dev.	Min	Max
Δ Share married	741	-0.0337	0.0327	-0.1415	0.1361
Δ Share 1^{st} married	741	-0.0118	0.0333	-0.1260	0.1365
Δ Share divorced / married	741	-0.0049	0.0455	-0.1932	0.1743
Δ Employment rate	741	0.0033	0.0330	-0.1434	0.1142
Δ Full-time rate, men	741	0.0181	0.0466	-0.2040	0.1545
Δ Full-time gap (M$-$F)	741	-0.0111	0.0466	-0.1501	0.1628
Δ Mean ln earnings, men	741	0.3395	0.1160	-0.0200	0.8091
Δ Mean ln earnings, women	741	0.3562	0.1074	0.0092	0.6557
Δ Earnings gap (M$-$F)	741	-0.0167	0.1231	-0.5763	0.4048
Δ S.d. of ln earnings	741	-0.0193	0.0835	-0.2488	0.2973
Δ Zero-earnings share	741	-0.0092	0.0341	-0.1402	0.1033
Δ AI-exposed employment share	741	0.0249	0.0299	-0.0858	0.1098

Table B2: Baseline CZ characteristics (2010): summary statistics

Variable (2010 CZ shares)	N	Mean	Std. dev.	Min	Max
Female share	741	0.493	0.017	0.430	0.524
Age 60+ share	741	0.098	0.017	0.054	0.172
College (BA+) share	741	0.198	0.061	0.090	0.449
High-AI occupations, women	741	0.147	0.030	0.037	0.247
High-AI occupations, men	741	0.280	0.051	0.162	0.466

Notes: Proportions of the 18–64 CZ population in 2010, from ACS microdata with survey weights allocated to commuting zones using PUMA–CZ crosswalk shares. "High-AI occupations" are those above the 66^{th} percentile of the Webb (2019) AI exposure distribution.

Table B3: Alternative aggregations of the ABS technology module, 2010–2021

	Δ Share married	Δ Share 1^{st} married	Δ Share divorced / married
Baseline (mean of five)	0.00857**	0.0130***	-0.0187***
	(0.00389)	(0.00391)	(0.00610)
Machine learning only	0.00682**	0.0104***	-0.0149***
	(0.00330)	(0.00340)	(0.00493)
NLP only	0.00903**	0.0137***	-0.0197***
	(0.00376)	(0.00371)	(0.00591)
Machine vision only	0.00686**	0.0104***	-0.0150***
	(0.00346)	(0.00362)	(0.00529)
Guided vehicles only	0.0109**	0.0165***	-0.0237***
	(0.00477)	(0.00470)	(0.00818)
Cognitive tools only	0.00935**	0.0142***	-0.0204***
	(0.00412)	(0.00411)	(0.00653)
Any-tool lower bound	0.00970**	0.0147***	-0.0212***
	(0.00450)	(0.00457)	(0.00723)

Each row rebuilds the shift-share on a different aggregation of the five ABS technology questions and re-estimates the fully controlled 2SLS specification. Coefficients are scaled by the standard deviation of their own exposure measure so that rows are comparable. The any-tool lower bound takes the maximum across the five tools, which is what the share of firms using at least one tool must exceed. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B4: Industries used in the AI shift–share exposure, with the adoption rate of each technology

NAICS	**Industry**	**ML**	**Vision**	**NLP**	**Voice**	**AGV**	**AIAdopt**
518	Data processing, hosting, and related services	11.3	3.7	7.6	3.5	0.4	5.30
5415	Computer systems design	7.8	2.6	4.7	4.2	0.4	3.94
511	Publishing	7.0	2.7	4.1	2.8	0.5	3.42
5417	Scientific research and development services	8.3	3.2	1.8	1.7	0.6	3.12
333	Machinery	7.5	5.6	0.7	0.8	0.6	3.04
334	Computer and electronic products	5.0	5.7	0.8	1.3	0.6	2.68
332	Fabricated metal products	6.6	4.1	0.8	0.8	0.8	2.62
322	Paper products	4.2	2.9	1.5	0.8	1.0	2.08
336	Transportation equipment	3.6	3.8	0.5	1.0	1.3	2.04
326	Rubber and plastic products	3.2	4.7	0.7	0.7	0.7	2.00
621–623	Health care	1.2	0.9	1.3	6.2	0.2	1.96
5411	Legal services	0.7	0.4	2.4	5.8	0.3	1.92
335	Electrical equipment, appliances, and components	3.9	4.0	0.5	0.5	0.6	1.90
5416	Management, scientific, and technical consulting services	2.5	0.8	1.7	2.5	0.2	1.54
5413	Architectural, engineering, and related services	2.2	2.3	0.7	1.7	0.7	1.52
5418	Advertising, public relations, and related services	3.1	0.5	1.4	2.3	0.2	1.50
5419	Other professional, scientific and technical services	2.2	1.1	1.2	2.7	0.3	1.50
517	Telecommunications	2.0	0.9	1.4	2.8	0.2	1.46
339	Miscellaneous manufacturing	4.0	1.8	0.5	0.7	0.2	1.44
337	Furniture and related products	3.7	1.7	0.5	0.7	0.3	1.38
331	Basic metals	2.9	3.0	0.2	0.2	0.2	1.30
321	Wood products	3.4	2.2	0.3	0.2	0.2	1.26
5412	Accounting, tax preparation, bookkeeping, and payroll services	2.4	1.0	0.9	1.6	0.2	1.22
52	Finance and insurance	1.5	0.6	0.9	2.2	0.6	1.16
325	Chemicals and chemical products	2.1	1.8	0.3	0.7	0.5	1.08
324	Coke and refined petroleum products	0.9	1.5	1.3	0.9	0.6	1.04
327	Nonmetallic mineral products	2.7	1.3	0.5	0.3	0.3	1.02
323	Printing and related support activities	2.8	1.2	0.5	0.4	0.2	1.02
61	Education	2.0	0.5	0.7	1.7	0.2	1.02
624	Social assistance	1.0	0.5	1.2	2.0	0.4	1.02
21	Mining, extraction, and support activities	2.2	0.8	0.3	1.1	0.3	0.94
5414	Specialized design services	0.9	0.6	0.8	1.9	0.4	0.92
311–312	Food, beverages, tobacco	1.6	1.3	0.3	0.8	0.5	0.91
313–316	Textile, apparel, leather products	2.2	0.9	0.2	0.8	0.3	0.88
81	Other services	1.8	0.8	0.5	0.7	0.5	0.86
53	Real estate	1.0	0.4	0.5	1.9	0.4	0.84
42	Wholesale trade	1.5	0.9	0.5	0.9	0.4	0.84
48–49	Transportation and storage	0.8	0.6	0.6	1.0	0.9	0.78
11	Agriculture, forestry, fisheries	1.1	1.3	0.2	0.0	1.2	0.76
44–45	Retail trade	1.2	0.7	0.5	0.6	0.5	0.70
55	Management of companies and enterprises	0.2	0.2	0.0	2.7	0.0	0.62
56	Administrative and support service	1.0	0.3	0.4	1.0	0.2	0.58
72	Accommodation and food services	0.9	0.5	0.6	0.6	0.3	0.58
23	Construction	0.9	0.4	0.4	0.8	0.3	0.56
22	Utilities	0.8	0.3	0.5	1.0	0.0	0.52
71	Arts, entertainment, and recreation	0.4	0.4	0.3	1.1	0.2	0.48

All 46 industry cells that carry a U.S. shifter, ordered by adoption. Columns 3–7 give the percentage of firms in the industry reporting active use of machine learning, machine vision, natural language processing, voice recognition and automated guided vehicles; within each technology the rate is firm-weighted across the ABS rows composing the industry. The final column, AIAdopt_j^{US}, is their unweighted mean and is the shifter used throughout. Source: ABS Technology Module 2020 (NSF 23-313, tables 40–43, 47).

Table B5: NACE–NAICS correspondence and European adoption rates used to build the instrument

NACE Rev. 2	**Matched NAICS**	**AIAdopt^{EU} (%)**
J62_J63	518	28.17
M72	5417	25.40
J61	517	19.83
M69–M71	5411–5416	18.08
J58–J60	511, 519	17.97
C19	324	17.81
C26	334, 339	16.83
C21	3254	15.73
M73–M75	5418–5419	12.50
C28	333	11.85
C29_C30	336	11.39
C27	335	10.39
C20	325	9.77
D_E	22	8.86
C22_C23	326–327	7.91
N	56	7.19
L68	53	7.11
G46	42	6.51
C16–C18	321–323	6.35
G47	44	6.00
C24_C25	331–332	5.94
H	48	5.22
C10–C12	311–312	5.04
C31–C33	337	4.70
F	23	4.62
C13–C15	313	3.66
I	72	3.27

The 27 NACE Rev. 2 cells for which Eurostat publishes sectoral AI adoption, the U.S. industries each is matched to, and the European adoption rate. The European measure is the share of enterprises with ten or more employees using at least one AI technology in 2021; Eurostat does not publish the sectoral breakdown by technology, so it is not constructed the same way as the U.S. shifter. The figure is Eurostat's own EU-27 aggregate and no further country weighting is applied. Source: Eurostat, ICT Usage in Enterprises, 2021.

Table B6: Robustness: Adao–Kolesar–Morales exposure-robust standard errors, 2010–2021

	Δ Share married			Δ Share 1^{st} married			Δ Share divorced / married		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	4.592***	3.789***	1.718***	3.542***	3.738***	2.087***	-5.171***	-5.348***	-4.888***
	(1.398)	(1.095)	(0.504)	(1.175)	(1.146)	(0.695)	(1.636)	(1.566)	(1.363)
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	5.572**	4.980**	3.786**	5.047**	5.962**	5.758**	-6.369**	-7.072**	-8.262**
	(2.551)	(2.200)	(1.712)	(2.372)	(2.707)	(2.676)	(2.925)	(3.125)	(3.581)

Standard errors are corrected for the shift-share structure following Adao, Kolesar and Morales (2019), which treats the industry shocks rather than the commuting zones as the independent units. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B7: Robustness: controlling for baseline telework capacity, 2010–2021

	Δ Share married			Δ Share 1^{st} married			Δ Share divorced / married		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	1.678*	1.443	1.404*	1.286	1.589	1.695*	-4.230***	-4.415***	-4.551***
	(0.918)	(0.928)	(0.809)	(1.105)	(1.007)	(0.931)	(1.262)	(1.431)	(1.432)
Telework capacity	0.119***	0.112***	0.168**	0.0923**	0.102**	0.210**	-0.0385	-0.0443	-0.180
	(0.0372)	(0.0409)	(0.0821)	(0.0405)	(0.0422)	(0.0837)	(0.0398)	(0.0392)	(0.115)
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	2.850**	2.586*	3.358**	4.112**	4.844***	5.256***	-6.878***	-7.307***	-7.837***
	(1.400)	(1.429)	(1.499)	(1.741)	(1.716)	(1.629)	(2.137)	(2.591)	(2.517)
Telework capacity	0.0943**	0.0905**	0.149*	0.0324	0.0423	0.175**	0.0176	0.00889	-0.148
	(0.0442)	(0.0451)	(0.0775)	(0.0554)	(0.0527)	(0.0853)	(0.0619)	(0.0552)	(0.107)

All specifications add the 2010 commuting-zone telework capacity index built from Dingel and Neiman (2020). All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B8: Robustness: controlling for the pre-period change in the outcome, 2010–2021

	Δ Share married			**Δ Share 1^{st} married**			**Δ Share divorced / married**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	5.311***	4.462***	1.472	4.247***	4.367***	2.053**	-7.613***	-7.324***	-4.709***
	(0.505)	(0.627)	(1.001)	(0.609)	(0.776)	(0.944)	(0.921)	(1.186)	(1.226)
Δ_{05}^{10} Share married	-0.368***	-0.371***	-0.419***						
	(0.0605)	(0.0624)	(0.0689)						
Δ_{08}^{10} Share 1^{st} married				-0.431***	-0.439***	-0.477***			
				(0.0528)	(0.0477)	(0.0455)			
Δ_{05}^{10} Share divorced / married							-0.515***	-0.535***	-0.595***
							(0.0597)	(0.0594)	(0.0624)
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	6.736***	6.203***	3.721**	6.270***	7.113***	6.256***	-9.913***	-10.02***	-8.409***
	(0.984)	(1.121)	(1.873)	(0.660)	(1.005)	(1.625)	(1.302)	(1.684)	(2.361)
Δ_{05}^{10} Share married	-0.388***	-0.389***	-0.414***						
	(0.0664)	(0.0673)	(0.0670)						
Δ_{08}^{10} Share 1^{st} married				-0.461***	-0.466***	-0.475***			
				(0.0556)	(0.0505)	(0.0487)			
Δ_{05}^{10} Share divorced / married							-0.558***	-0.567***	-0.593***
							(0.0692)	(0.0669)	(0.0626)

Each column controls for the pre-period change in its own outcome: 2005–2010 for the share married and the divorced-to-married ratio, 2008–2010 for the first-marriage share, since the ACS marital-history questions begin in 2008. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B9: Robustness: accounting for robot exposure, 2010–2021

	Δ Share married			Δ Share 1^{st} married			Δ Share divorced / married		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	4.829***	4.084***	2.068**	3.632***	3.866***	2.237**	-5.449***	-5.717***	-5.505***
	(0.441)	(0.606)	(0.990)	(0.570)	(0.731)	(1.052)	(1.068)	(1.334)	(1.469)
Robot exposure	0.00159***	0.00153***	0.00107**	0.000605	0.000667	0.000460	-0.00186**	-0.00192**	-0.00189**
	(0.000527)	(0.000538)	(0.000510)	(0.000498)	(0.000563)	(0.000590)	(0.000799)	(0.000892)	(0.000855)
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	5.906***	5.489***	4.564**	5.201***	6.246***	6.298***	-6.762***	-7.718***	-9.613***
	(0.709)	(0.916)	(1.881)	(0.632)	(1.003)	(1.894)	(1.176)	(1.847)	(2.689)
Robot exposure	0.00172***	0.00170***	0.00138**	0.000792	0.000945	0.000958	-0.00202**	-0.00215**	-0.00240***
	(0.000554)	(0.000558)	(0.000572)	(0.000514)	(0.000577)	(0.000612)	(0.000813)	(0.000912)	(0.000889)

All specifications add the 2010–2019 change in local robot exposure constructed following Acemoglu and Restrepo (2020). All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B10: Robustness: baseline wealth and pre-period economic trajectory, 2010–2021

	Δ Share married			Δ Share 1^{st} married			Δ Share divorced / married		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	2.270**	0.924	0.875	1.161	1.119	1.095	-3.442***	-3.420***	-3.560***
	(0.935)	(1.010)	(0.953)	(1.085)	(1.166)	(1.149)	(1.104)	(1.250)	(1.278)
Median HH income	7.47e-07***	8.62e-07***	7.08e-07	7.54e-07**	7.88e-07**	7.34e-07	-5.34e-07*	-5.80e-07**	-9.75e-07
	(2.69e-07)	(2.78e-07)	(4.27e-07)	(3.19e-07)	(3.27e-07)	(4.45e-07)	(2.94e-07)	(2.77e-07)	(5.82e-07)
Income growth	0.0910	0.117	0.0430	-0.0646	-0.0825	-0.0753	0.236	0.263	0.110
	(0.166)	(0.178)	(0.218)	(0.180)	(0.175)	(0.192)	(0.344)	(0.355)	(0.397)
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	3.863**	2.193	3.030*	3.831**	4.838**	5.100**	-5.433***	-6.003**	-6.810***
	(1.848)	(1.607)	(1.781)	(1.806)	(1.958)	(2.002)	(1.952)	(2.411)	(2.225)
Median HH income	4.91e-07	6.79e-07*	4.54e-07	3.26e-07	2.52e-07	2.61e-07	-2.14e-07	-2.08e-07	-5.91e-07
	(4.05e-07)	(3.47e-07)	(4.23e-07)	(4.57e-07)	(4.26e-07)	(4.68e-07)	(4.64e-07)	(3.94e-07)	(6.02e-07)
Income growth	0.100	0.117	0.0604	-0.0492	-0.0812	-0.0430	0.224	0.262	0.0842
	(0.177)	(0.177)	(0.215)	(0.186)	(0.173)	(0.188)	(0.347)	(0.350)	(0.376)

All specifications add year-2000 median household income and 2000–2010 income growth from Chetty et al. (2014). All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B11: Robustness: predetermined 2003 industry shares, 2010–2021 outcomes

	Δ Share married			**Δ Share 1^{st} married**			**Δ Share divorced / married**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: OLS Estimates									
AI Exposure	4.081***	3.005***	0.807	3.351***	3.400***	1.673**	-4.307***	-4.088***	-2.827**
	(0.594)	(0.718)	(0.649)	(0.590)	(0.651)	(0.744)	(0.985)	(1.165)	(1.158)
Panel B: 2SLS Estimates									
AI Exposure^{2sls}	5.910***	5.236***	3.667**	5.352***	6.269***	5.578***	-6.755***	-7.436***	-8.004***
	(0.693)	(1.061)	(1.786)	(0.555)	(1.023)	(1.812)	(1.136)	(1.992)	(2.987)

Exposure is rebuilt on 2003 rather than 2010 commuting-zone industry employment shares. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B12: An occupation-based exposure measure, alone and against the industry shift–share

	Occupation exposure, alone	AI Exposure, joint	Occupation exposure, joint
Δ Share married	-0.278**	3.535**	-0.367**
	(0.108)	(1.436)	(0.176)
Δ Share 1^{st} married	-0.132	5.416***	-0.500***
	(0.107)	(1.504)	(0.146)
Δ Share divorced / married	0.162	-7.917***	0.505**
	(0.185)	(2.481)	(0.216)
Δ Employment rate	0.0503	9.528***	0.425*
	(0.149)	(2.366)	(0.221)
Δ Earnings gap (M−F)	0.0611	19.84***	0.266
	(0.347)	(5.346)	(0.543)
Δ S.d. of ln earnings	-0.719**	4.898	0.985*
	(0.311)	(4.574)	(0.511)

Occupation-based exposure is the 2010 employment-weighted mean of the Webb (2019) occupational percentile in the commuting zone, on a 0–1 scale. It measures how AI-exposed the local workforce already was, rather than how much the local industry mix went on to adopt, and the two are close to orthogonal across commuting zones (correlation -0.198). Column 1 estimates it alone by OLS, dropping the baseline high-AI occupational shares from the control set because they are constructed from the same index. Columns 2 and 3 enter it alongside the instrumented shift-share. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Table B13: Descriptive statistics: changes in 1^{st} marriage by couple type (2010–2021)

Variable ($\Delta_{2010-2021}$)	N	Mean	Std. dev.	Min	Max
Δ Share high–high 1^{st} marriages	741	0.0060	0.0112	-0.0347	0.0539
Δ Share mixed 1^{st} marriages	741	-0.0050	0.0141	-0.0499	0.0428
Δ Share low–low 1^{st} marriages	741	-0.0071	0.0103	-0.0340	0.0309

B.1 The shocks assumption

This appendix reports the diagnostics behind the identification discussion of Section 4.2. Identification in the exogenous-shifts framework (Borusyak et al., 2023; Adão et al., 2019) requires that industry-level shocks be as good as randomly assigned conditional on controls, and that there be enough of them. Neither condition is comfortable here, and the diagnostics say so directly. The design draws on 46 industries, but aggregate employment weights are concentrated: the effective number of shocks, computed as the inverse Herfindahl of those weights, is 13.5, and the five largest industries carry 0.52 of the total. The Rotemberg weights (Goldsmith-Pinkham et al., 2020) are more concentrated still, corresponding to an effective 8.0 industries, with the five largest carrying 0.58 of the identifying variation.

The shock-permutation test sharpens the point. When the adoption vector is reshuffled at random across the 46 industries and the model re-estimated two thousand times, the true industry ranking is not extreme: 0.74 of reshuffles produce a larger standardised coefficient for the share married, 0.84 for first marriage and 0.84 for the divorce ratio. The reason is visible in the placebo exposures themselves, which correlate 0.85 with the true exposure on average. The commuting-zone share matrix has few effective dimensions, so almost any assignment of these 46 numbers to industries generates a similar geographic map. The same fact explains why every alternative aggregation of the ABS technologies in Appendix Table B3 yields the same answer.

Table B14 pursues this further. Column 2 adds the share-weighted commuting-zone projection of seven baseline industry characteristics (college share, mean log earnings, female share, mean age, mean Webb score, full-time share and log employment), which is the approach Borusyak et al. (2023) recommend when shocks are suspected of correlating with observable industry traits; the first-marriage estimate survives while the other two lose precision. Column 3 goes further and purges the adoption vector itself of those characteristics before rebuilding the shift-share. Those characteristics account for 0.34 of the variance of industry adoption, and on the residual the estimates are indistinguishable from zero. No single industry is responsible for the result: when each is dropped in turn, the fully controlled 2SLS coefficient for the share married ranges from 3.578 to 4.597, with $|t|$ never falling below

2.14.

Table B14: Identification: conditioning the shift-share on industry structure

	Baseline	+ industry traits	Residualised shifter
Δ Share married	3.786**	1.538	-1.439
	(1.720)	(4.940)	(2.892)
Δ Share 1^{st} married	5.758***	11.54**	-1.715
	(1.729)	(5.637)	(3.051)
Δ Share divorced / married	-8.262***	-15.13	-1.235
	(2.695)	(10.26)	(3.292)
Shock-permutation test			
Δ Share married: share of reshuffles more extreme	0.74		
Δ Share 1^{st} married: share of reshuffles more extreme	0.84		
Δ Share divorced / married: share of reshuffles more extreme	0.84		

Column 1 is the baseline 2SLS estimate. Column 2 adds the share-weighted commuting-zone projection of seven 2010 industry characteristics (college share, mean log earnings, female share, mean age, mean Webb AI-exposure score, full-time share and log employment), which is what Borusyak, Hull and Jaravel (2022) recommend when shocks may correlate with observable industry traits. Column 3 instead purges the adoption vector itself of those characteristics and rebuilds the shift-share on the residual, estimated by OLS. All controls are measured at the base year of the window. Models are weighted by base-year CZ population. Standard errors in parentheses are clustered by state. $***p < 0.01, **p < 0.05, *p < 0.1$.

Taken together, these diagnostics bound the interpretation without overturning it. The estimand is the effect of specialising in the industries that went on to adopt AI, relative to other commuting zones over the same period, and carries no claim about the effect of AI holding constant everything correlated with adoption. That estimand answers the question posed in the main text, whether the AI wave affected places the way the robot wave did, and it is the object Section 5 reports.

C Supplementary figures

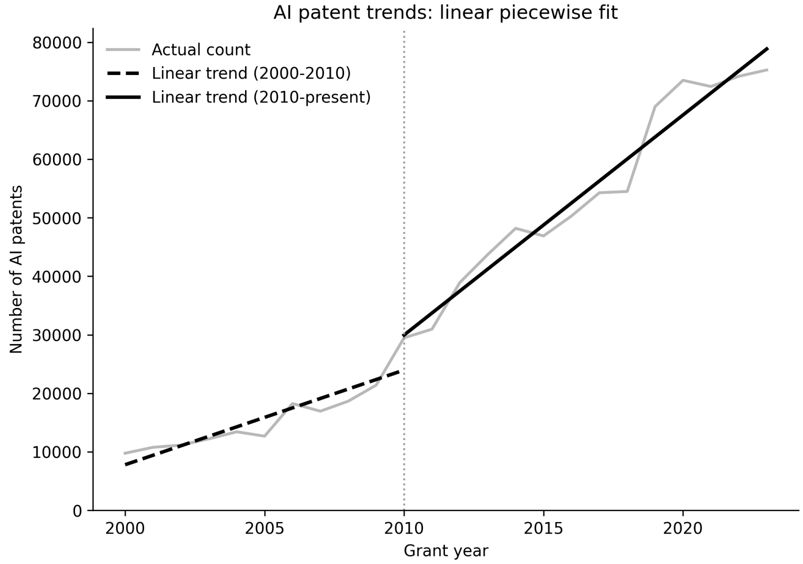

Figure C1: **AI patent trends: piecewise linear fit.** Annual counts of AI patents granted by the USPTO, with separate linear trends fitted before and after 2010.

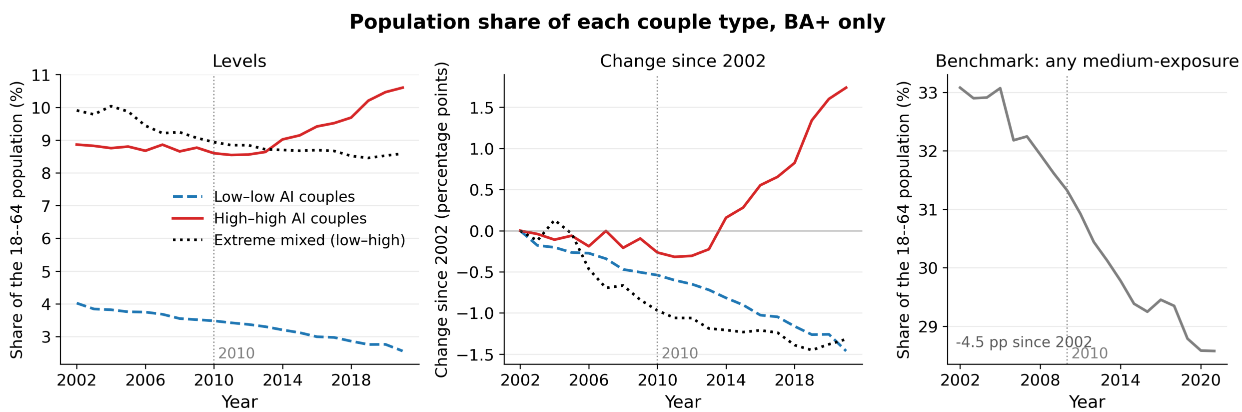

Figure C2: **Population share of each couple type, college graduates only.** This figure replicates Figure 2 panel by panel, restricting the entire universe — both the couples in the numerator and the total 18–64 household population in the denominator — to individuals holding a Bachelor's degree or higher. The continued divergence between high–high and low–low couples indicates that the baseline trends are not driven solely by the educational wage premium.