Consent and Careers: Legal Protection and Women's Occupational Opportunities in Malawi

Alexander Nketiah
University of Alicante

September 2026

Abstract

This study examines whether changes in early sexual debut associated with Malawi's 2011 age-of-consent reform are linked to women's occupational outcomes in early adulthood. Using five rounds of the Malawi Demographic and Health Survey from 2000 to 2024, I combine birth-cohort exposure to the reform with district-by-ethnicity variation in the historical prevalence of sexual debut before age 16. The interaction between cohort exposure and historical intensity is used as an instrument for early sexual debut. The first stage shows that exposed cohorts in historically higher-intensity communities have lower rates of sexual debut before age 16. The same variation is associated with a higher probability of skilled occupation, a higher occupational-status score, and a higher probability of professional-or-clerical employment. The reduced-form estimates provide the clearest evidence of this occupational pattern. The 2SLS estimates point in the same direction, but their large magnitudes and a significant first-stage pre-exposure pattern make the IV interpretation more dependent on the identifying assumptions. Overall, the findings link reform-associated variation in early sexual debut to improved occupational attainment in early adulthood.

Keywords: Age-of-consent laws; Early sexual debut; Women's occupational attainment; Instrumental variables; Difference-in-differences; Malawi.

1 Introduction

Early sexual initiation, child marriage, and teenage childbearing can constrain women's human-capital accumulation and later economic opportunities. These transitions are closely connected during adolescence: early pregnancy and early marriage can interrupt schooling, reduce mobility, and shift household responsibilities, while school dropout can increase exposure to cohabitation, marriage, and childbearing. These constraints are particularly important in developing countries, where gender inequalities in education, autonomy, and employment remain substantial (Jayachandran, 2015). Evidence from Malawi likewise documents close relationships between schooling, sexual debut, pregnancy, and marriage during adolescence (Glynn et al., 2010, 2018; Sunny et al., 2019; Wado et al., 2019). Less is known, however, about whether legal interventions that alter these early-life transitions are reflected in women's later economic opportunities.

This paper studies that question using Malawi's 2011 Penal Code amendment, which raised the statutory age of sexual consent from 13 to 16 (Government of Malawi, 2011; Kangaude and Skelton, 2018). Previous evidence for Malawi shows that the reform is associated with delayed sexual debut, lower teenage childbearing and early marriage, and improvements in schooling-related outcomes (Nketiah, 2026). I extend this analysis to a later stage of the life course and ask whether changes in early sexual debut associated with the reform-intensity design are also associated with women's occupational opportunities in adulthood.

There are several reasons why delaying early sexual debut could matter for later occupation. Sexual debut occurs at an age when decisions about schooling, marriage, and fertility can have persistent consequences for human-capital accumulation. Schooling expansions can improve later labor-market outcomes (Duflo, 2001), while interventions that increase girls' schooling can reduce teenage childbearing and other early-life transitions (Black et al., 2008; Duflo et al., 2015). Evidence from Malawi also shows that adolescent interventions can jointly affect schooling, marriage, and pregnancy (Baird et al., 2011). Early marriage and teenage childbearing have been linked to subsequent educational and labor-market outcomes (Field and Ambrus, 2008; Soler-Hampejsek et al., 2021; Wilson, 2022; Ashcraft et al., 2013; Fletcher and Wolfe, 2009; Levine, 2014). These findings suggest that a reform that changes sexual behavior during adolescence may have consequences extending beyond immediate demographic outcomes and into women's later occupational attainment.

Estimating this relationship is difficult because early sexual debut is not randomly assigned. Women who initiate sex before age 16 may differ systematically in household resources, schooling opportunities, local labor markets, marriage norms, and other characteristics that also influence later occupation. I address this endogeneity by combining

cohort exposure to the 2011 reform with district-by-ethnicity variation in the historical prevalence of sex before age 16. The identifying idea is that the increase in the age of consent had greater scope to alter behavior in communities where sexual debut before 16 had historically been more common. The interaction between reform exposure and historical early-debut intensity is therefore used as the excluded instrument for early sexual debut. Cohort-by-intensity diagnostics and robustness exercises assess whether this variation is associated with pre-existing cohort patterns.

The analysis pools five rounds of the Malawi Demographic and Health Survey: 2000, 2004, 2010, 2015–16, and 2024. The main specification compares women born in 1999–2006, who were young enough to be exposed to the higher consent threshold during adolescence, with women born in 1987–1994, who had largely completed the relevant age window before the reform. Cohorts born in 1995–1998 are excluded because their exposure to the reform is ambiguous.[1] Historical treatment intensity is constructed from older cohorts who completed adolescence before the reform, ensuring that the estimation sample does not mechanically determine the intensity measure. In the analytical sample, mean historical intensity is similar for control and exposed cohorts, which provides a descriptive check that the intensity measure is not concentrated mechanically in one cohort group.

The empirical results show two main patterns. First, the reform-intensity interaction significantly predicts lower sexual debut before age 16 among exposed cohorts in historically higher-intensity communities. Second, the same reform-intensity variation is positively associated with skilled occupation, higher occupational-status scores, and professional-or-clerical employment. The reduced-form estimates provide the most direct occupational evidence. OLS and 2SLS estimates are directionally consistent, with early sexual debut associated with poorer occupational outcomes, but the 2SLS estimates are large relative to outcome means and rely on stronger identifying assumptions. I therefore interpret their magnitudes cautiously.

This study contributes to two related literatures. First, it extends evidence on adolescent legal protection and related early-life transitions beyond immediate demographic and schooling outcomes to women's occupational opportunities (Baird et al., 2011; Field and Ambrus, 2008; Duflo et al., 2015; Nketiah, 2026). This matters because the economic relevance of adolescent legal protection beyond adolescence depends partly on whether behavioral changes during adolescence are reflected in later economic outcomes. Second, it connects the literature on age-of-consent and statutory-rape laws with research on women's labor-market opportunities. Existing evidence on these laws has concentrated primarily on adolescent fertility, particularly in higher-income settings (Frakes and Hard-

[1] The 2015–16 survey is treated as one survey round. The exact rounds are stated for reproducibility because the occupation analysis depends on observing exposed cohorts after they enter adulthood.

ing, 2015; Lepage, 2022). By examining occupational outcomes in Malawi, this paper provides evidence from a low-income setting where enforcement and access-to-justice constraints may make the economic consequences of adolescent transitions especially important.

The interpretation remains deliberately bounded by the research design. The instrument generates variation in early sexual debut but does not separately identify the roles of schooling, marriage, and fertility in transmitting that variation to later occupational outcomes. These factors are therefore treated as plausible intermediate pathways rather than as independently identified causal mechanisms.

The remainder of the paper proceeds as follows. Section 2 describes Malawi's legal reform and institutional context. Section 3 presents the data, sample construction, and outcome measurement. Section 4 develops the empirical strategy and identifying assumptions. Section 5 presents the main results and cohort-by-intensity diagnostics. Section 6 reports robustness and validity checks. Section 7 discusses limitations, and Section 8 concludes.

2 Institutional Background

2.1 Malawi's 2011 Age-of-Consent Reform

Before 2011, Malawi's statutory age of sexual consent was 13. The 2011 Penal Code amendment raised the threshold to 16, extending legal protection against sexual activity involving girls below that age (Government of Malawi, 2011; Kangaude and Skelton, 2018). The reform therefore created differential exposure across birth cohorts: girls who had not yet reached age 16 when the reform took effect could be protected by the new threshold during adolescence, whereas older cohorts had already passed the relevant age window.

The reform established a clear statutory threshold at age 16 without a close-in-age exemption. Its potential influence therefore extends beyond formal prosecution. By increasing the legal consequences associated with sexual relationships involving girls below 16, the reform may alter the incentives of potential partners while also providing parents, schools, and community actors with a stronger legal basis for discouraging such relationships.

The same legal threshold could also generate partner-market spillovers. If some potential partners respond to the reform by avoiding relationships with girls below 16, relationship demand may shift toward slightly older adolescents or young women. Such substitution would mean that some women above the threshold could be indirectly affected by the reform even if they are not themselves below age 16. This possibility is one reason I avoid interpreting the estimates as a pure direct effect of the statutory threshold alone and

instead focus on the reform-intensity variation defined by cohort exposure and historical early-debut prevalence.

Whether these statutory changes translated into behavior depends on the institutional context. Enforcement of sexual-offense laws in Malawi is constrained by limited institutional capacity and by difficulties in reporting and prosecuting offenses occurring within private relationships and household networks (Kangaude, 2017, 2021). A UNICEF study of survivors of sexual violence in Malawi likewise documents barriers to accessing police and court services (UNICEF Malawi, 2019). The reform may therefore operate not only through the direct deterrent effect of expected legal sanctions but also through changes in social expectations and community responses. A statutory threshold can provide a focal point for acceptable behavior and strengthen the position of families and other local actors seeking to discourage early sexual relationships (Sunstein, 1996; Acemoglu and Jackson, 2017). The empirical analysis does not attempt to distinguish these mechanisms. Instead, it examines whether behavioral responses to the reform are subsequently reflected in women's occupational outcomes.

2.2 Demographic, Social, and Labor-Market Context

Malawi provides an important setting in which to examine these long-run outcomes. The economy is predominantly rural, informal employment is widespread, agriculture accounts for a substantial share of employment, and gender disparities in paid work and earnings persist (Mussa, 2016; UN Women, 2024; World Bank, 2021). In this context, occupational category provides a useful, although incomplete, measure of women's economic opportunities. Movement from agricultural, domestic, or unskilled work toward skilled manual, service, clerical, technical, managerial, or professional employment can represent economically meaningful occupational upgrading even when conventional measures of earnings or formal employment are unavailable.

Adolescent sexual behavior may affect later occupational attainment through several interconnected pathways. One important pathway is human-capital accumulation. Early sexual initiation can increase exposure to pregnancy, marriage, and related household responsibilities during the years in which girls would otherwise be completing secondary education or acquiring skills. Interruptions at this stage may reduce subsequent access to occupations that require greater schooling, training, or specialization. Evidence from Malawi shows that school failure, sexual debut, pregnancy, and marriage are closely related during adolescence (Glynn et al., 2018; Sunny et al., 2019). A systematic review of causal evidence from low- and middle-income countries finds heterogeneous and generally modest effects of education on sexual and reproductive-health outcomes (Psaki et al., 2019).

A second pathway operates through the timing of family formation. Delaying sexual debut may postpone pregnancy and marriage, allowing young women to enter adulthood with fewer immediate childcare and household responsibilities. This may increase the time and flexibility available for continued education, labor-market entry, occupational search, or movement into work outside agriculture and domestic activities. These considerations are particularly relevant in a labor market in which women's occupational opportunities may depend not only on whether they work, but also on the type of work they can enter.

These pathways provide an economic rationale for examining occupational attainment as a longer-run outcome of adolescent legal protection. They should not, however, be interpreted as separately identified mechanisms in the empirical analysis. Schooling, marriage, and childbearing may occur jointly and may themselves respond to the reform. The instrumental-variables design therefore estimates the occupational relationship associated with instrument-induced variation in early sexual debut but does not separately identify the contribution of any one intermediate pathway.

A further identification concern is the Marriage, Divorce and Family Relations Act of 2015, which established 18 as the minimum legal age for marriage (Government of Malawi, 2015). Because this reform occurred after the 2011 age-of-consent amendment and could independently affect marriage, schooling, fertility, and subsequent labor-market outcomes, changes observed among younger cohorts cannot automatically be attributed to a single institutional pathway. Malawi also amended the Penal Code again in 2023, introducing a general definition of a child as a person under 18 and amending provisions on sexual offences involving children, including a close-in-age defence in the "sexual activity with a child" provision (Government of Malawi, 2023). This later amendment does not change the first-stage outcome of sexual debut before age 16, since the youngest 2006 birth cohort had reached age 16 by 2022, but it may affect the later schooling, relationship, or early labor-market environment observed for the youngest cohorts in the 2024 DHS. These overlapping reforms are particularly important for interpreting mechanisms. Accordingly, the analysis focuses on whether the variation generated by exposure to the 2011 reform and historical early-debut intensity predicts occupational outcomes, rather than attempting to decompose the estimated effect into separate schooling, marriage, or fertility channels.

Figure 1 summarizes this framework. The 2011 reform, interacted with historical early-debut intensity, provides the source of variation in early sexual debut, while education, marriage, and childbearing represent potential intermediate pathways to later occupational outcomes. These pathways motivate the analysis but are not separately identified by the instrumental-variables design.

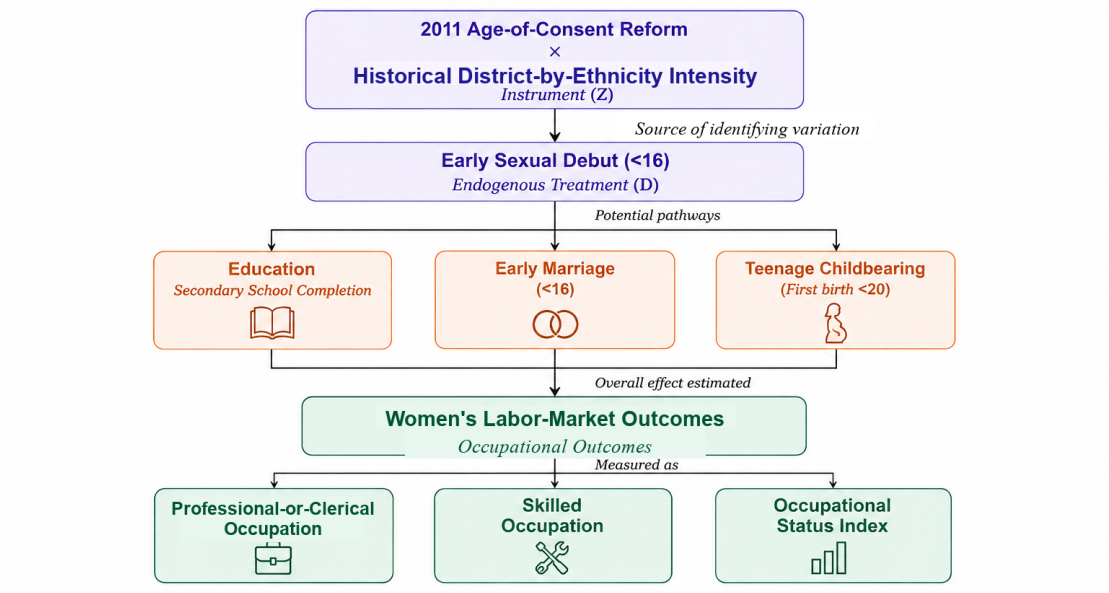

Figure 1: Conceptual framework

3 Data

3.1 Survey Data

The empirical analysis uses five rounds of the Malawi Demographic and Health Survey (DHS): 2000, 2004, 2010, 2015–16, and 2024 (National Statistical Office (NSO) [Malawi] and ORC Macro, 2001, 2005; National Statistical Office (NSO) [Malawi] and ICF Macro, 2011; National Statistical Office (NSO) [Malawi] and ICF, 2017; National Statistical Office (NSO) [Malawi] et al., 2026). The DHS provides nationally representative repeated cross-sectional data on women of reproductive age, including information on demographic histories, education, household characteristics, religion, residence, and current occupation. The surveys also contain retrospective information on age at first sexual intercourse, age at first marriage, and birth histories, which makes it possible to compare birth cohorts that reached adolescence before and after the 2011 reform and to examine their occupational outcomes later in life.

Because the DHS consists of repeated cross-sections rather than a longitudinal panel, the same women are not followed over time. The analysis therefore does not compare changes in an individual woman's occupation before and after the reform. Instead, it compares women from differently exposed birth cohorts across district-by-ethnicity groups with different historical levels of early sexual debut. Variables are harmonized across survey rounds using definitions that can be constructed consistently, with particular attention to sexual debut, birth timing, district, ethnicity, and occupation. DHS sampling weights are used throughout the analysis to account for differences in survey design and sampling probabilities across rounds.

Occupational information is harmonized into three measures: skilled occupation, an occupational-status index, and professional-or-clerical occupation. Their construction and interpretation are described in detail in Section 3.4. Appendix Table A1 reports the definitions and observed universe of the treatment and outcome variables, while Appendix Table A7 provides the occupational crosswalk used to construct the three measures.

3.2 Analytical Sample

The main analytical sample is defined using birth cohorts to distinguish women according to their exposure to the 2011 reform. The control group consists of women born between 1987 and 1994, who had reached or passed age 16 by the time the statutory threshold changed. The exposed group consists of women born between 1999 and 2006, who were young enough for the higher age-of-consent threshold to apply during adolescence. Cohorts born between 1995 and 1998 fall between these groups because their exposure around the relevant ages is less clear. I therefore exclude these intermediate cohorts from the main specification rather than assigning them to either the control or exposed group. Appendix Figure A1 summarizes the cohort assignment.

The occupational analysis is further restricted to women aged 18 or older at the time of interview. This restriction ensures that occupational outcomes are observed only after respondents have entered adulthood. It does not, however, eliminate differences in labor-market maturity between the cohort groups. Exposed women are younger on average than control women, which is particularly relevant for entry into skilled and professional-or-clerical occupations. The main specification therefore includes birth-year and survey-year fixed effects, and Section 6.2 examines robustness to restricting the sample to women aged 22 or older.

Applying the age-18 restriction yields 20,047 occupation-eligible women: 14,235 from the control cohorts and 5,812 from the exposed cohorts. Requiring a classifiable occupational category further reduces the sample to 19,307 women, comprising 13,868 controls and 5,439 exposed women. The 740 age-eligible observations that are excluded at this stage cannot be assigned to the harmonized occupational measures. The resulting 19,307 women constitute the main occupational estimation sample used for the reduced-form, OLS, first-stage, and 2SLS specifications. Appendix Table A8 summarizes the sample construction.

The pooled data contain two post-reform DHS rounds, 2015–16 and 2024. Because the occupational analysis requires respondents to be at least 18 and the oldest exposed cohort was born in 1999, exposed women do not enter the occupational estimation sample until the 2024 survey. The 2015–16 round nevertheless contributes observations from the control cohorts. Survey-year fixed effects absorb shocks common to respondents observed

in the same DHS round, while the reform-intensity interaction continues to vary across birth cohorts and district-by-ethnicity cells.

For comparability across the empirical analysis, the first-stage specification is estimated on the same 19,307 women used for the occupational reduced-form and 2SLS estimates. Using a common estimation sample ensures that differences across the first-stage, reduced-form, and IV results are not generated by changes in sample composition.

Table 1 reports weighted descriptive statistics for the control and exposed cohorts satisfying the occupational-analysis requirements. The exposed cohorts are younger by construction, with a mean age of 21.2 compared with 25.4 among control women. They are also somewhat more rural and have a slightly lower mean wealth quintile. Historical early-debut intensity is similar across the two groups, averaging 38.1 percent among controls and 39.1 percent among exposed women. These statistics describe the composition of the pooled analytical sample and should not be interpreted as evidence of balance arising from randomized treatment assignment.

Table 1: Descriptive Balance by Cohort Group

	Control cohorts			Exposed cohorts		
	Mean	SD	Count	Mean	SD	Count
Age	25.389	5.235	14,235	21.210	2.223	5,812
Rural residence	0.809	0.393	14,235	0.842	0.365	5,812
Wealth quintile	3.047	1.444	14,235	2.887	1.460	5,812
Catholic	0.170	0.375	14,235	0.135	0.341	5,812
Protestant/Christian	0.588	0.492	14,235	0.367	0.482	5,812
Muslim	0.196	0.397	14,235	0.330	0.470	5,812
No religion	0.044	0.206	14,235	0.167	0.373	5,812
Other religion	0.002	0.042	14,235	0.001	0.030	5,812
Pre-reform intensity (%)	38.149	11.113	14,235	39.111	11.606	5,812
First sex before age 16	0.329	0.470	14,235	0.266	0.442	5,812
Professional-or-clerical occupation	0.060	0.238	13,868	0.067	0.250	5,439
Skilled occupation	0.104	0.305	13,868	0.111	0.314	5,439
Occupational status index, 0–5	1.100	1.229	13,868	0.939	1.334	5,439
Age-18-plus eligible observations		14,235			5,812	
Main occupational estimation observations		13,868			5,439	

Notes: DHS sampling weights are used. Counts are variable-specific. The age-18-plus eligible sample contains 20,047 observations. The main occupational estimation sample contains 19,307 observations.

3.3 Treatment and Historical Intensity

The endogenous treatment is an indicator equal to one if a woman reports having had first sexual intercourse before age 16. The age-16 threshold corresponds directly to the statutory threshold introduced by the 2011 reform. Because age at first sexual intercourse

is reported retrospectively in the DHS, the measure may be subject to recall or reporting error.

The second component of the empirical design is historical early-debut intensity, defined as the historical prevalence of sexual debut before age 16 within a respondent's district-by-ethnicity cell. The measure is constructed using women born in 1970 or earlier. These women completed adolescence well before the 2011 reform and are distinct from the birth cohorts used in the main exposed-versus-control comparison. Constructing intensity from these older cohorts therefore avoids mechanically defining a community's historical exposure using the same women whose occupational outcomes enter the main analysis.

The district-by-ethnicity grouping allows historical early-debut prevalence to vary both across districts and across ethnic groups within districts. Ethnicity adds a relatively persistent dimension of community background to the geographic grouping, although it does not eliminate concerns related to migration or imperfect measurement of respondents' historical location. The resulting measure should therefore be interpreted as a proxy for historical community exposure rather than as a direct measure of enforcement intensity or current sexual behavior.

The main analysis contains 72 district-by-ethnicity cells across 27 districts. Fourteen self-reported ethnicity categories are observed in the analysis sample, with Chewa, Lomwe, Yao, Ngoni, and Tumbuka accounting for the largest shares. The number of retained ethnicity cells per district ranges from one to seven, with a median of three. Cells are retained only when at least 25 valid observations are available to construct the historical share, limiting the influence of highly imprecise cell-level estimates.

The identifying variation therefore does not come solely from differences across districts or solely from comparisons between ethnic groups within the same district. Instead, the empirical design compares exposed-versus-control cohort differences across district-by-ethnicity cells with different historical levels of early sexual debut. District-by-ethnicity fixed effects absorb permanent level differences across these cells, while identification comes from how cohort differences vary with the predetermined historical-intensity measure.

Because historical intensity is constructed from older cohorts and then assigned to later respondents according to their district-by-ethnicity cell, it is predetermined relative to the cohorts entering the main reform-exposure comparison. Predeterminedness does not imply that historical intensity is mechanically unrelated to birth cohort in the pooled DHS data. Some cohorts may be represented differently across cells. Table 1 shows that mean intensity is nevertheless similar between the control and exposed cohorts, while Appendix Table A3 reports intensity means by individual birth year. The empirical design therefore does not rely on exact raw balance between cohort and intensity; instead, it relies on the

fixed-effects structure and the cohort-by-intensity diagnostics described in Section 4.

Historical early-debut intensity varies substantially across the 72 cells. The cell-level mean is approximately 38.8 percent, with a standard deviation of 11.8 percentage points and a range of approximately 11.0 to 66.5 percent. Appendix Table A2 reports the corresponding cell-level summary statistics, and Appendix Figure A2 shows the distribution of the standardized measure. Before estimation, the continuous intensity measure is standardized to have mean zero and unit standard deviation. Accordingly, a one-unit increase in standardized intensity corresponds to an approximately 11.8 percentage-point increase in the underlying historical prevalence of sexual debut before age 16.

There is also substantial descriptive variation across major ethnicity groups. Mean historical intensity is approximately 50 percent among Yao and Lomwe respondents, compared with about 32 percent among Chewa and Ngoni respondents and 29 percent among Tumbuka respondents. These averages are descriptive and should not be interpreted as the source of identification, which operates at the district-by-ethnicity cell level. They nevertheless illustrate the extent of variation in the historical measure.

The continuous standardized measure is used in the main specification because it retains the full variation in historical intensity across district-by-ethnicity cells. As a robustness check, Section 6.2 instead classifies cells as above or below the median historical intensity to assess whether the results are sensitive to using a continuous measure.

3.4 Occupational Outcome Measurement

The analysis focuses on occupational attainment rather than employment alone. This distinction is important in Malawi, where employment is highly heterogeneous and a large share of economic activity occurs in agriculture and informal work. The outcomes are therefore designed to capture differences in the type and relative status of work women perform rather than simply whether they report being employed.

The first primary outcome, *skilled occupation*, is a binary indicator distinguishing relatively skilled work from agricultural, domestic, sales, and unskilled manual activities. It equals one for women classified in services or skilled manual work, clerical work, or professional, technical, and managerial occupations. The measure is intended to capture whether a woman has moved into occupational categories requiring relatively greater skills or specialization.

The second primary outcome is the *occupational status index*. This measure preserves more information about the occupational distribution by ranking women across six ordered categories. The index ranges from zero for women not currently working to five for professional, technical, or managerial employment, with intermediate categories rep-

resenting agricultural or unskilled manual work, sales or domestic work, services or skilled manual work, and clerical employment. A one-unit change in this index therefore represents movement between adjacent occupational-status categories rather than a percentage-point change.[2]

The third outcome, *professional-or-clerical occupation*, is a binary indicator for clerical, professional, technical, or managerial work. [3] I treat this measure as secondary rather than co-primary because entry into these occupations generally requires more schooling and labor-market experience, while the reform-exposed cohorts are still relatively young. The outcome is nevertheless informative about whether the reform-intensity pattern extends to a more selective occupational margin.

Finally, the analysis distinguishes occupational type from selection into employment. Women who are not currently working occupy the lowest category of the occupational status index, while the skilled and professional-or-clerical indicators describe occupational position within the relevant observed universe. Because differences in employment itself could influence the composition of women for whom occupational categories are observed, the robustness analysis separately examines whether the reform-intensity interaction predicts current employment and re-estimates the occupational specifications among women who are working. Appendix Table A7 reports the complete mapping from DHS occupational categories into each constructed outcome.

4 Empirical Design

4.1 Treatment Intensity and Cohort Exposure

The empirical strategy combines cohort exposure to Malawi's 2011 age-of-consent reform with cross-sectional variation in the historical prevalence of sexual debut before age 16. Let $Exposed_c$ indicate whether birth cohort c was young enough to be exposed to the higher statutory threshold during adolescence, as defined in Section 3.2, and let $Intensity_g$ denote the standardized historical prevalence of sex before age 16 in district-by-ethnicity cell g. The interaction

[2]The occupational status index is an ordinal measure. Its coefficients should consequently be interpreted as movements in the constructed occupational ranking rather than as proportional changes in earnings, productivity, or job quality.

[3]The professional-or-clerical indicator is an occupational classification rather than a measure of formal-sector employment. The DHS occupation categories identify the type of work reported by respondents but do not establish whether a formal contract covers employment, provides social protection, or otherwise satisfies a definition of formal employment. The indicator should therefore be interpreted as capturing entry into professional, technical, managerial, or clerical occupations rather than formal-sector employment per se.

$$Z_{cg} = Exposed_c \times Intensity_g$$

is used as the excluded instrument for sexual debut before age 16.

The empirical strategy is related to cohort-by-intensity designs that combine differential exposure across birth cohorts with cross-sectional variation in either policy intensity or pre-intervention conditions. Duflo (2001), for example, combines cohort exposure to Indonesia's school-construction program with regional differences in the intensity of school construction. Related studies of malaria eradication combine cohort exposure to national eradication campaigns with pre-existing geographic variation in malaria intensity (Bleakley, 2010; Lucas, 2010; Cutler et al., 2010). In the present setting, however, $Intensity_g$ is not a measure of policy implementation or enforcement intensity. It is a predetermined measure of the historical prevalence of early sexual debut. The interaction therefore captures whether the difference between exposed and control cohorts varies systematically with the extent to which sex before age 16 had historically been common in a community.

The underlying hypothesis is that raising the statutory age of consent from 13 to 16 had greater scope to alter early sexual behavior in communities where sexual debut before age 16 was historically more prevalent. Because $Intensity_g$ is standardized, coefficients on the reform-intensity interaction represent the differential change associated with a one-standard-deviation increase in historical early-debut intensity.

The unit g is a district-by-ethnicity cell. The identifying comparison therefore does not rely solely on differences across districts or solely on comparisons between ethnic groups within the same district. Instead, it compares exposed-versus-control cohort differences across district-by-ethnicity cells with different historical early-debut intensities. Variation in intensity occurs both across districts and across ethnic groups within districts, so both dimensions contribute to the identifying variation.

District-by-ethnicity fixed effects absorb permanent level differences across cells, birth-year fixed effects absorb national differences across cohorts, and survey-year fixed effects absorb shocks common to respondents observed in the same DHS round. Identification therefore does not come from persistent level differences between high- and low-intensity communities or from national cohort changes shared across all communities. Rather, it comes from whether the difference between exposed and control cohorts varies systematically with predetermined historical early-debut intensity.

The key identifying requirement is therefore a parallel-trends-type assumption for the intensity gradient. In the absence of the 2011 reform, differences between exposed and control cohorts in early sexual debut and occupational outcomes should not have evolved systematically with historical early-debut intensity. A violation would arise, for example, if higher-intensity communities were already experiencing different cohort-specific changes

in early sexual behavior or occupational opportunities that would have continued even without the reform. The cohort-by-intensity diagnostics presented below are used to assess the plausibility of this assumption.

4.2 First Stage and Reduced Form

The first stage examines whether the reform-intensity interaction predicts sexual debut before age 16:

$$D_{icgt} = \pi \left(Exposed_c \times Intensity_g \right) + X'_{icgt}\Gamma + \lambda_g + \delta_c + \tau_t + u_{icgt}, \tag{1}$$

where D_{icgt} equals one if woman i, born in cohort c and assigned to district-by-ethnicity cell g, reports having had first sexual intercourse before age 16. The vector X_{icgt} contains the individual-level controls included in the main specification: religion and rural residence. The terms λ_g, δ_c, and τ_t denote district-by-ethnicity, birth-year, and survey-year fixed effects, respectively.

The coefficient of interest, π, captures how the exposed-versus-control cohort difference in early sexual debut varies with historical early-debut intensity. Because $Intensity_g$ is standardized, π represents the differential change associated with a one-standard-deviation increase in historical intensity. A negative coefficient indicates that exposed cohorts experience a larger reduction in the probability of sexual debut before age 16 in historically higher-intensity communities. The first stage therefore provides evidence on the relevance of the reform-intensity interaction as an instrument for early sexual debut.

Current household wealth is not included in the main control vector because it is measured at the time of the survey and may itself reflect earlier schooling, marriage, fertility, and labor-market trajectories. Controlling for it could therefore condition on a variable that is partly determined by post-adolescent outcomes. As a robustness check, I add household wealth to assess whether the estimates are sensitive to its inclusion.

I next estimate the reduced-form relationship between the reform-intensity interaction and each occupational outcome:

$$Y_{icgt} = \rho \left(Exposed_c \times Intensity_g \right) + X'_{icgt}\Phi + \lambda_g + \delta_c + \tau_t + e_{icgt}, \tag{2}$$

where Y_{icgt} denotes skilled occupation, the occupational status index, or professional-or-clerical occupation. The coefficient ρ captures how the exposed-versus-control cohort difference in each occupational outcome varies with historical early-debut intensity.

The reduced form is informative because it relates the reform-intensity variation directly

to occupational outcomes without rescaling the estimates by the first-stage coefficient. Presenting the first stage and reduced form separately therefore addresses two distinct empirical questions: whether the reform-intensity interaction predicts early sexual debut and whether the same source of variation is associated with differential occupational outcomes among exposed cohorts. The 2SLS specification in the following section then links these two relationships under the additional instrumental-variables assumptions.

4.3 Instrumental-Variables Estimation

The second-stage equation is

$$Y_{icgt} = \beta \widehat{D}_{icgt} + X'_{icgt}\Theta + \lambda_g + \delta_c + \tau_t + \varepsilon_{icgt}, \tag{3}$$

where $\widehat{D}_{icgt}$ is the component of sexual debut before age 16 predicted by the reform-intensity interaction in the first stage. The coefficient β captures the relationship between instrument-induced variation in early sexual debut and the occupational outcome Y_{icgt}.

Because $D_{icgt} = 1$ denotes sexual debut before age 16, a negative value of β implies that instrument-induced increases in early sexual debut are associated with lower occupational attainment. Equivalently, variation in the instrument that reduces the probability of sexual debut before age 16 is associated with improved occupational outcomes.

Under the instrumental-variables assumptions discussed in Section 4.4, including relevance, the exclusion restriction, and monotonicity, the 2SLS coefficient can be given a local causal interpretation for women whose early-sexual-debut behavior responds to the reform-intensity instrument (Imbens and Angrist, 1994). Because the excluded instrument is continuous, this interpretation should not be understood as the effect for a single, directly identifiable group of compliers. Rather, the estimate reflects a weighted average of causal responses generated by changes in the instrument. It should therefore not be interpreted as a population-average effect for all Malawian women.

For comparison, I also estimate OLS regressions relating early sexual debut directly to each occupational outcome using the same control structure. These estimates are descriptive rather than causal because early sexual debut may be correlated with observed and unobserved determinants of later occupational attainment. The OLS estimates nevertheless provide a useful benchmark for comparing the direction and magnitude of the conditional association with the corresponding 2SLS estimates. Differences between OLS and 2SLS should not be interpreted as discrepancies between two estimates of the same population parameter, since the IV estimator relies on instrument-induced variation and may capture effects at a different margin of the population.

4.4 Identification Assumptions

A causal interpretation of the 2SLS estimates requires several assumptions. First, instrument relevance requires the reform-intensity interaction to predict sexual debut before age 16. This condition is assessed using the first-stage coefficient together with the weak-instrument diagnostics reported below. The interaction must generate sufficient variation in early sexual debut for the second-stage relationship to be identified.

Second, the cohort-by-intensity design requires a parallel-trends-type assumption. In the absence of the 2011 reform, differences between exposed and control cohorts in early sexual debut and occupational outcomes should not have evolved systematically with historical early-debut intensity. This requirement is related to recent work on difference-in-differences designs with continuous treatment intensity, which emphasizes that identification and interpretation depend on how parallel-trends assumptions are formulated across different levels of treatment intensity (Callaway et al., 2024). Their framework does not directly characterize the instrumental-variables design used here, but it highlights the additional care required when identifying variation is continuous rather than binary.

District-by-ethnicity fixed effects absorb permanent differences across cells, while birth-year fixed effects absorb national differences across cohorts and survey-year fixed effects absorb shocks common to respondents observed in the same DHS round. These fixed effects do not, however, eliminate differential cohort trends associated with historical intensity. I therefore use cohort-by-intensity profiles to examine whether such differential patterns are evident among cohorts whose adolescence preceded the reform. The minimum-age robustness check addresses a separate concern by assessing whether differences in labor-market maturity between exposed and control cohorts account for the occupational results.

Third, the exclusion restriction requires the reform-intensity interaction to affect occupational outcomes through its effect on early sexual debut and the downstream consequences of that change, rather than through alternative pathways operating independently of a woman's early sexual debut. This assumption cannot be tested directly and is demanding in this setting. For example, the 2011 reform could affect schooling, marriage, fertility, parental supervision, or community behavior even among girls whose timing of sexual debut does not change. Such responses would create a direct pathway from the reform-intensity interaction to later occupational outcomes that does not operate through the endogenous treatment measured here.

Other institutional and programmatic changes could create similar concerns. The 2015 Marriage, Divorce and Family Relations Act may have had stronger effects in communities where early marriage and early sexual debut were historically more common. Schooling,

health, HIV-prevention, NGO, or girls' empowerment programs may also have been disproportionately targeted toward historically high-intensity communities and could influence later occupational outcomes independently of sexual debut. In addition, the partner-market responses discussed in Section 2 may affect relationship opportunities in ways that are not fully mediated by a woman's own timing of sexual debut. The cohort-by-intensity diagnostics and robustness exercises provide indirect evidence on these possibilities, but they cannot establish the exclusion restriction.

Fourth, a local IV interpretation requires monotonicity. Movements in the reform-intensity instrument should shift early sexual behavior in the same direction among women whose behavior responds to the instrument. In particular, stronger reform-intensity exposure should not cause some women to initiate sex before age 16 when they would otherwise have delayed it while simultaneously causing others to delay early sexual debut. Under monotonicity and the other IV assumptions, the 2SLS coefficient can be interpreted as a weighted average of local causal responses among women whose early-sexual-debut behavior is affected by the instrument (Imbens and Angrist, 1994). Because the instrument is continuous, this interpretation does not correspond to a single, directly identifiable group of compliers.

Finally, the model contains one excluded instrument and one endogenous regressor and is therefore exactly identified. Consequently, no overidentifying-restrictions test is available for assessing the exclusion restriction. Schooling, marriage, and childbearing may all form part of the pathway linking instrument-induced changes in early sexual debut to later occupational outcomes, but the design does not provide separate exogenous variation in these intermediate outcomes. The 2SLS coefficient therefore incorporates any downstream pathways operating through changes in early sexual debut and should not be interpreted as identifying the contribution of any one mechanism.

4.5 Inference and Diagnostic Tests

All regressions use DHS sampling weights. Standard errors are clustered at the district-by-ethnicity level, which corresponds to the level at which historical early-debut intensity varies. The main estimation sample contains 72 district-by-ethnicity clusters. Because cluster-robust inference can be sensitive to the number and distribution of clusters, the main 2SLS results are supplemented with wild-cluster-bootstrap inference (Cameron et al., 2008; Roodman et al., 2019).

Instrument strength is evaluated using the first-stage coefficient together with the Kleibergen-Paap rk Wald weak-identification F statistic for the excluded instrument. These diagnostics are considered jointly rather than treating statistical significance of the first stage alone as sufficient evidence of a well-behaved IV specification (Keane and Neal, 2024).

In particular, the strength of the first stage is interpreted alongside the magnitude of the reduced-form relationship and the resulting 2SLS estimates. A statistically relevant first stage does not, by itself, establish the validity of the instrument or imply that the resulting IV magnitude is economically plausible.

The empirical design is further assessed using several complementary validity and robustness checks. Cohort-by-intensity profiles examine whether differential relationships with historical intensity are already visible among cohorts whose adolescence preceded the reform. A minimum-age restriction addresses differences in labor-market maturity between exposed and control cohorts, while alternative specifications assess sensitivity to the inclusion of household wealth, the definition of historical intensity, employment selection, and geographic composition. These exercises cannot establish the identifying assumptions, but they provide evidence on whether the main results depend on particular features of the specification or sample.

To examine differential cohort patterns more directly, I replace the single $Exposed_c \times Intensity_g$ interaction with interactions between standardized historical intensity and individual birth-year indicators:

$$Y_{icgt} = \sum_{k \neq 1991} \theta_k \left(\mathbf{1}\{c = k\} \times Intensity_g \right) + X'_{icgt}\Psi + \lambda_g + \delta_c + \tau_t + \varepsilon_{icgt}. \tag{4}$$

The 1991 birth cohort is omitted, so each θ_k measures the historical-intensity gradient for cohort k relative to the corresponding gradient for the 1991 cohort. These estimates are used as cohort-by-intensity diagnostic profiles, rather than as a conventional staggered-treatment event study. The pre-exposure coefficients are examined individually and jointly to assess whether the historical-intensity gradient was already changing across cohorts before exposure to the 2011 reform.

5 Main Results

5.1 First Stage

Table 2 reports the first-stage relationship between the reform-intensity interaction and sexual debut before age 16. To maintain a common estimation sample across the first-stage, reduced-form, and 2SLS specifications, the first stage is estimated on the same 19,307 women included in the main occupational estimation sample. This ensures that differences across the three stages of the analysis are not driven by changes in sample composition.

Table 2: First-Stage Estimate

	First sex before age 16
Exposed cohort x standardized intensity	-0.0442***
	(0.00922)
Observations	19,307
Clusters	72
KP weak-ID F	23.09
Excluded-instrument p-value	0.0000

Notes: Standard error is in parentheses. District-by-ethnicity, birth-year, and survey-year fixed effects are included. Donut cohorts are excluded. DHS sampling weights are used. Standard error is clustered by district-by-ethnicity cell. KP weak-ID F is the Kleibergen-Paap rk Wald weak-identification F statistic. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The estimated coefficient on $Exposed_c \times Intensity_g$ is -0.0442 (SE $= 0.0092$). Because historical early-debut intensity is standardized, this coefficient implies that a one-standard-deviation increase in historical intensity is associated with a 4.42 percentage-point more negative exposed-versus-control cohort difference in the probability of sexual debut before age 16. In the underlying intensity measure, one standard deviation corresponds to approximately 11.8 percentage points. Thus, moving from a community with historical early-debut prevalence near the sample mean of 38.8 percent to one approximately one standard deviation above that mean is associated with a 4.42 percentage-point larger reduction in early sexual debut among exposed cohorts relative to control cohorts.

The magnitude is economically meaningful relative to the control-cohort mean of approximately 32.9 percent: 4.42 percentage points corresponds to roughly 13 percent of that baseline prevalence. The Kleibergen-Paap rk Wald weak-identification F statistic is 23.09, and the excluded-instrument coefficient is statistically different from zero at conventional levels. Taken together, these results show that the reform-intensity interaction provides statistically relevant variation in early sexual debut in the main estimation sample. Importantly, however, first-stage relevance does not by itself establish that this variation can be interpreted as exogenous or attributed entirely to the 2011 reform.

Figure 2 provides the corresponding cohort-by-intensity diagnostic profile for sexual debut before age 16. Each coefficient measures the historical-intensity gradient for a given birth cohort relative to the omitted 1991 cohort. The coefficients for many of the later cohorts are negative, indicating that the relationship between historical intensity and early sexual debut becomes more negative for these cohorts relative to 1991. This pattern is directionally consistent with the negative first-stage estimate.

However, the diagnostic does not show a clean pattern emerging only among reform-exposed cohorts. The joint test of the pre-exposure cohort-by-intensity coefficients yields a p-value of 0.0111, rejecting the null of no differential pre-exposure pattern. In other

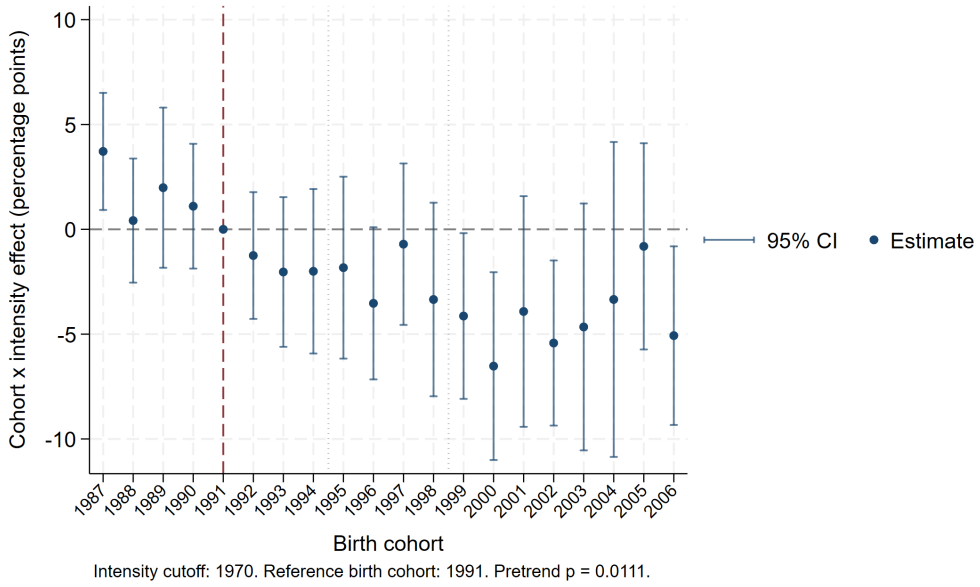

Figure 2: Cohort-by-intensity diagnostic profile for sexual debut before age 16.

Notes: Coefficients are birth-cohort interactions with standardized historical early-debut intensity relative to the omitted 1991 birth cohort. Coefficients are in percentage points. Vertical bars report 95% confidence intervals. Joint pretrend test $p = 0.0111$.

words, early sexual debut was already evolving differently across the historical-intensity distribution among cohorts whose adolescence preceded the reform.

The significant pre-exposure pattern complicates a simple causal interpretation of the first-stage estimate. The negative coefficient in Table 2 establishes that the instrument predicts early sexual debut, but the pre-exposure pattern makes it difficult to interpret the entire exposed-cohort shift as originating with the 2011 reform. The first stage therefore provides evidence of instrument relevance, but weaker evidence for the stronger claim that the reform generated a new, otherwise absent change in early sexual behavior across the intensity distribution. This distinction is important for the 2SLS interpretation because the IV estimates require not only relevance but also the identifying assumptions discussed in Section 4.4.

The first-stage pretrend does not automatically imply that the occupational reduced-form relationships are spurious. Those outcomes have separate cohort-by-intensity diagnostics and must be evaluated on their own pre-exposure patterns. Nevertheless, because the endogenous treatment itself exhibits a statistically detectable pre-exposure pattern, I interpret the 2SLS estimates more cautiously and place greater emphasis on the occupational reduced-form evidence, which directly relates the reform-intensity interaction to later occupational outcomes without rescaling by the first-stage coefficient.

5.2 Reduced Form, OLS, and 2SLS Estimates

Table 3 reports the reduced-form, OLS, and 2SLS estimates for the three occupational outcomes.

Table 3: Reduced-Form, OLS, and 2SLS Estimates for Occupational Outcomes

	Skilled	Status index	Professional-or-clerical
Panel A. Reduced form			
Exposed cohort x standardized intensity	0.0297***	0.135***	0.0250***
	(0.00902)	(0.0370)	(0.00697)
Panel B. OLS associations			
Early sexual debut	-0.0199***	-0.0533**	-0.0199***
	(0.00579)	(0.0223)	(0.00435)
Panel C. 2SLS estimates			
Early sexual debut (instrumented)	-0.671***	-3.043***	-0.566***
	(0.190)	(0.776)	(0.166)
Observations	19,307	19,307	19,307
Outcome mean	0.106	1.054	0.062
KP weak-ID F	23.09	23.09	23.09
Wild-bootstrap p-value	0.003	0.002	0.000

Notes: Standard errors are in parentheses. District-by-ethnicity, birth-year, and survey-year fixed effects are included. Donut cohorts are excluded. DHS sampling weights are used. Standard errors are clustered by district-by-ethnicity cell. Binary outcomes are in probability units; occupational status is a 0–5 index. KP weak-ID F is the Kleibergen-Paap rk Wald weak-identification F statistic for Panel C. Wild-bootstrap p-values refer to Panel C and use 999 replications and Webb weights. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel A presents the reduced-form estimates. A one-standard-deviation increase in historical early-debut intensity is associated with a 2.97 percentage-point larger exposed-versus-control difference in the probability of skilled occupation (SE = 0.0090). For the occupational-status index, the corresponding difference is 0.135 index units (SE = 0.0370), while for professional-or-clerical occupation it is 2.50 percentage points (SE = 0.0070). All three estimates are statistically significant at the 1 percent level. The corresponding outcome means are 10.6 percent for skilled occupation, 1.054 for the occupational-status index, and 6.2 percent for professional-or-clerical occupation.

The positive reduced-form coefficients indicate that occupational outcomes are relatively better for exposed cohorts in communities with higher historical early-debut intensity. These estimates are particularly informative because they relate the reform-intensity interaction directly to later occupational outcomes without converting that relationship into an estimated effect of early sexual debut. They therefore provide the most direct evidence on the occupational pattern generated by the cohort-by-intensity design. Their

causal interpretation nevertheless continues to depend on the identifying assumption that, in the absence of the reform, occupational outcomes would not have evolved differentially across cohorts according to historical intensity.

Panel B reports the OLS associations between sexual debut before age 16 and occupational outcomes. Early sexual debut is associated with a 1.99 percentage-point lower probability of skilled occupation, a 0.053-unit lower occupational-status score, and a 1.99 percentage-point lower probability of professional-or-clerical occupation. These coefficients show a consistently negative conditional relationship between early sexual debut and later occupational attainment. They should not be interpreted as causal effects, however, because women who experience sexual debut before age 16 may differ from other women in observed and unobserved characteristics that also influence their occupational outcomes.

Panel C reports the 2SLS estimates obtained by instrumenting sexual debut before age 16 with the reform-intensity interaction. The estimated coefficients are negative for all three outcomes: -0.671 for skilled occupation, -3.043 for the occupational-status index, and -0.566 for professional-or-clerical occupation. The negative signs are consistent with the first-stage and reduced-form results: the instrument is negatively associated with early sexual debut and positively associated with occupational outcomes. Under the IV assumptions discussed in Section 4.4, this pattern implies that instrument-induced reductions in early sexual debut are associated with better occupational outcomes.

The magnitude of the 2SLS coefficients, however, requires caution. For skilled occupation, the reduced-form coefficient of 0.0297 is rescaled by the first-stage coefficient of approximately -0.0442, producing an IV estimate close to -0.67. The resulting estimate is large relative to the skilled-occupation mean of 0.106. Similarly, the professional-or-clerical coefficient of -0.566 is large relative to an outcome mean of 0.062, while the -3.043 estimate for the occupational-status index is large relative to its mean of 1.054 on the 0–5 scale. The large magnitudes therefore partly reflect the combination of sizeable occupational reduced forms and a more modest first-stage change in early sexual debut.

For the two binary outcomes, the 2SLS coefficients are linear probability coefficients and are not constrained to remain within the unit interval. Because the binary outcomes are estimated using linear 2SLS, the coefficients represent IV effects in probability units rather than literal population-wide changes in predicted probabilities. More generally, the difference between the OLS and 2SLS magnitudes does not imply that one estimator is correcting the other toward the same population effect. OLS relies on observed variation in early sexual debut, whereas 2SLS relies on the variation predicted by the instrument and, under the required assumptions, has a local rather than population-average interpretation.

The Kleibergen-Paap rk Wald weak-identification F statistic is 23.09. Wild-cluster-bootstrap p-values are 0.003 for skilled occupation, 0.002 for the occupational-status index, and below 0.001 for professional-or-clerical occupation, using 999 replications and Webb weights. These results indicate that the statistical significance of the IV estimates is robust to the alternative inference procedure. They do not, however, resolve the identifying concerns discussed in Section 5.1 and Section 4.4.

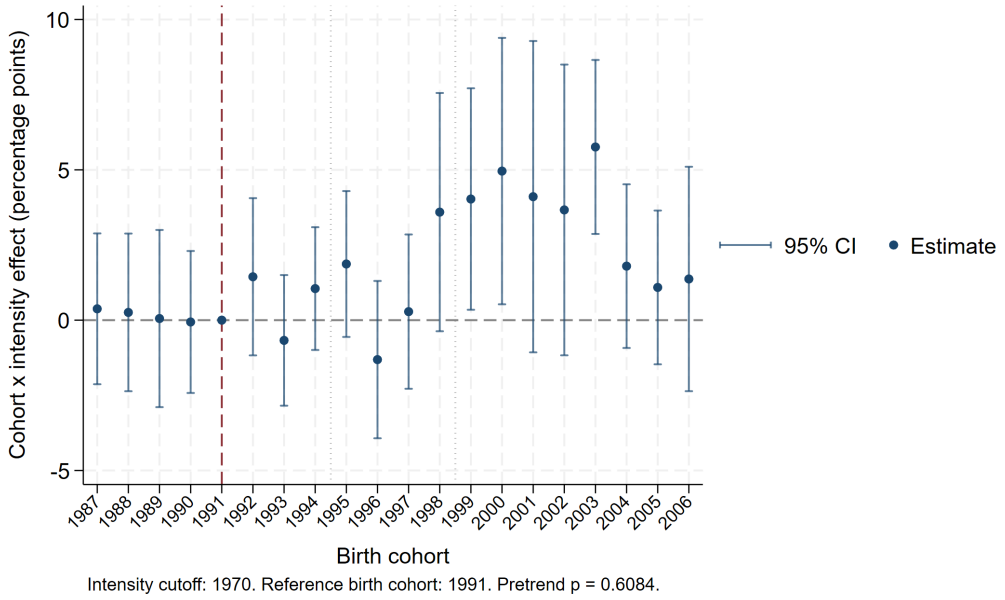

(a) Skilled occupation

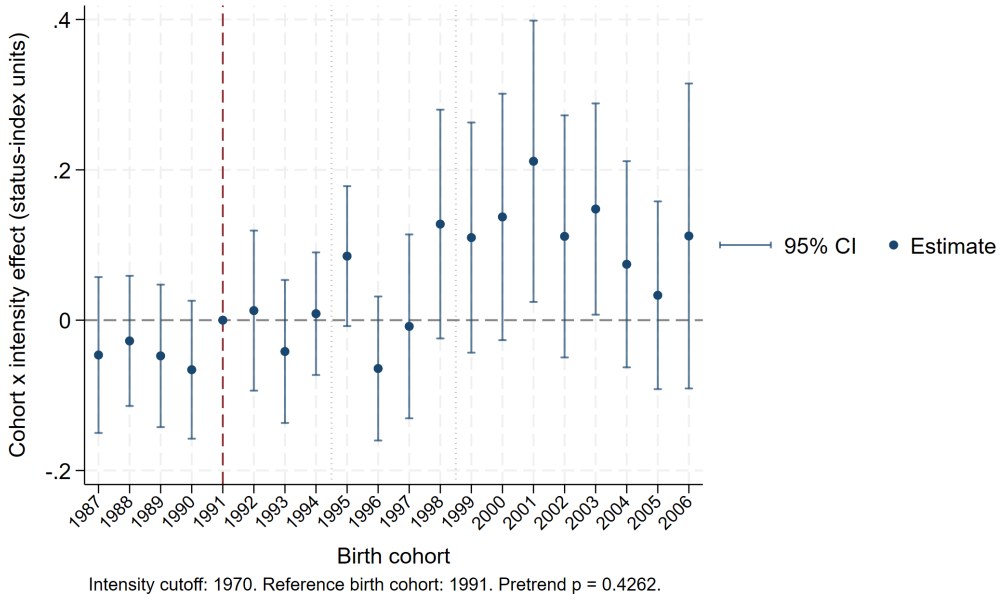

(b) Occupational status index

Figure 3: Cohort-by-intensity diagnostic profiles for primary occupational outcomes.

Figure 3 presents the cohort-by-intensity diagnostic profiles for skilled occupation and the

occupational-status index. For skilled occupation, the pre-exposure coefficients do not display a systematic pattern, and the joint test of the pre-exposure coefficients yields a p-value of 0.6084. The null of no differential pre-exposure pattern is therefore not rejected. The coefficients for many later cohorts are positive, although individual estimates are imprecisely estimated.

The occupational-status index displays a similar pattern. Its joint pre-exposure test yields a p-value of 0.4262, again providing no statistically detectable evidence of differential pre-exposure cohort patterns across the historical-intensity distribution. The professional-or-clerical profile, reported in Appendix Figure A3, produces a corresponding joint pre-exposure p-value of 0.5858. Thus, unlike the first-stage outcome, none of the three occupational outcomes shows a statistically detectable pre-exposure pattern in these joint tests.

Taken together, the occupational results show a consistent reduced-form pattern: exposed-versus-control differences in occupational attainment become more positive as historical early-debut intensity increases, and the occupational cohort profiles do not show statistically detectable differential pre-exposure patterns. The 2SLS estimates point in the corresponding direction when early sexual debut is instrumented with the reform-intensity interaction, but their large magnitudes and the pre-exposure pattern identified in the first stage make the IV interpretation more assumption-dependent. For this reason, the reduced-form estimates provide the central evidence on occupational outcomes, while the 2SLS estimates provide complementary evidence on the relationship with instrument-induced variation in early sexual debut.

The robustness analyses that follow examine whether the occupational patterns are sensitive to labor-market maturity, alternative control sets, treatment-intensity definitions, employment selection, and geographic composition.

6 Robustness and Validity Checks

6.1 Employment Selection

One concern is that the occupational results may reflect changes in whether women work rather than changes in the types of occupations they enter. This distinction is particularly relevant for the occupational-status index because women who are not currently working are assigned to the lowest category of the index. If the reform-intensity interaction mainly affected employment participation, a positive coefficient on occupational status could therefore arise even without movement toward higher-ranked occupations among working women.

Appendix Table A5 examines this possibility in two steps. Panel A uses current employment as the dependent variable. The coefficient on the reform-intensity interaction is 0.0135 (SE = 0.0260) and is not statistically significant. Thus, there is no statistically detectable relationship between the reform-intensity interaction and the probability of being currently employed in the age-eligible sample.

Panel B then re-estimates the occupational reduced forms among women who are currently working and have a classifiable occupation. The coefficients remain positive across all three outcomes: 0.0415 for skilled occupation, 0.162 for the occupational-status index, and 0.0414 for professional-or-clerical occupation. The first two are statistically significant at the 5 percent level and the professional-or-clerical estimate at the 1 percent level. These results show that the positive occupational pattern is also present when the comparison is restricted to women already participating in work.

Panel C reports the corresponding 2SLS estimates among currently working women. The coefficients remain negative: −0.775 for skilled occupation, −3.026 for the occupational-status index, and −0.773 for professional-or-clerical occupation. The Kleibergen-Paap rk Wald weak-identification F statistic is 24.18 in this restricted sample. Although the estimates are less precisely estimated than in the main sample, their directions are consistent with the main 2SLS results.

Taken together, these exercises indicate that the main occupational results are not explained solely by differences in employment participation. In particular, the persistence of the reduced-form relationships among currently working women suggests that the cohort-by-intensity pattern also appears in occupational composition, not only in whether women work.

The estimates conditional on current employment are nevertheless best viewed as a robustness check rather than as an alternative main specification. Employment may itself be influenced by earlier schooling, marriage, fertility, or other post-adolescent outcomes, so conditioning on current employment changes the composition of the estimation sample. The main specification therefore remains preferable for estimating the overall occupational relationship, while the worker-only results help distinguish occupational upgrading from simple changes in employment participation.

6.2 Alternative Specifications and Age Composition

Appendix Table A6 examines whether the occupational results are sensitive to alternative controls, the definition of historical intensity, age composition, and the inclusion of major urban centers. Panel A reports the reduced-form estimates and Panel B reports the corresponding 2SLS estimates. Across all four specifications, the reduced-form coefficients

remain positive and the 2SLS coefficients remain negative for each occupational outcome.

First, adding current household wealth produces estimates that are nearly unchanged from the main specification. The reduced-form coefficients are 0.0295 for skilled occupation, 0.134 for the occupational-status index, and 0.0249 for professional-or-clerical occupation. The corresponding 2SLS estimates are −0.674, −3.065, and −0.570. This similarity indicates that the main results are not sensitive to the exclusion of current household wealth from the baseline control set.

Second, I replace the continuous standardized intensity measure with an indicator for whether a district-by-ethnicity cell has historical early-debut intensity above the sample median. The reduced-form estimates remain positive and statistically significant, while the corresponding 2SLS estimates remain negative. Because the binary measure and the standardized continuous measure have different units, their coefficient magnitudes are not directly comparable. The relevant finding is that the direction of the results is preserved under an alternative way of classifying historical exposure.

Third, I restrict the sample to women aged 22 or older. This check addresses the concern that exposed women are younger on average and may therefore be observed at an earlier stage of occupational development. In the restricted sample of 12,336 women, the reduced-form coefficients remain positive: 0.0380 for skilled occupation, 0.168 for occupational status, and 0.0322 for professional-or-clerical occupation. The corresponding 2SLS coefficients remain negative at −0.572, −2.521, and −0.484, while the Kleibergen-Paap rk Wald weak-identification F statistic is 27.89. These results indicate that the main occupational pattern is not confined to women observed immediately after entering adulthood.

Finally, I exclude respondents from major urban centers to assess whether the results are driven by locations with substantially different labor-market opportunities. The reduced-form coefficients remain positive for all three occupational outcomes, and the 2SLS coefficients remain negative. The resulting sample contains 17,745 women, and the Kleibergen-Paap rk Wald weak-identification F statistic is 23.05. This exercise indicates that the main pattern is also present outside the major urban centers.

Overall, these specifications show that the direction of the occupational results is preserved when the control set, historical-intensity definition, age composition, and geographic composition of the sample are altered. Appendix Figures A4–A6 summarize the corresponding 2SLS robustness estimates graphically. These exercises support the stability of the empirical pattern across alternative specifications, while leaving the identifying assumptions discussed in Section 4.4 unchanged.

7 Limitations

The analysis is subject to some limitations. First, the first-stage cohort profile shows a statistically significant pre-exposure pattern. Early sexual debut was already evolving differently across the historical-intensity distribution before the reform, which makes it difficult to attribute the entire first-stage relationship to the 2011 reform. This is the main reason the 2SLS estimates are interpreted more cautiously than the occupational reduced-form results.

Second, the exclusion restriction cannot be tested directly. Other changes affecting younger cohorts may also have varied with historical early-debut intensity and influenced later occupational outcomes independently of early sexual debut. The 2015 Marriage, Divorce and Family Relations Act and the 2023 Penal Code amendment are particularly relevant overlapping institutional changes.

Third, the exposed cohorts are still relatively young when their occupational outcomes are observed. The analysis therefore captures early-adult occupational attainment rather than completed career trajectories. Moreover, the DHS occupational measures distinguish broad types of work but do not measure earnings, job stability, productivity, or long-run career progression.

8 Conclusion

This study examines whether changes in adolescent sexual behavior associated with Malawi's 2011 age-of-consent reform are linked to women's occupational opportunities in early adulthood. The empirical strategy combines cohort exposure to the higher statutory threshold with district-by-ethnicity variation in the historical prevalence of sexual debut before age 16.

The results show that exposed cohorts in historically higher-intensity communities have lower rates of sexual debut before age 16 and better occupational outcomes. In particular, the reform-intensity interaction is associated with a higher probability of skilled occupation, a higher occupational-status score, and a higher probability of professional-or-clerical employment. The reduced-form estimates provide the most direct evidence of this occupational pattern, while the 2SLS estimates link the same variation to instrument-induced changes in early sexual debut. The large IV magnitudes and the significant first-stage pre-exposure pattern, however, make the 2SLS interpretation more dependent on the identifying assumptions discussed in the paper.

Overall, the findings indicate that the cohort-by-intensity variation associated with lower early sexual debut is also associated with improved occupational attainment among young

women in Malawi. The analysis does not separately identify the roles of schooling, marriage, and fertility, and the exposed cohorts are still relatively early in their working lives. Future research that follows these cohorts further into adulthood could examine whether the occupational differences persist and clarify the pathways through which adolescent legal protection may affect later economic outcomes.

References

Acemoglu, D. and Jackson, M. O. (2017). Social norms and the enforcement of laws. *Journal of the European Economic Association*, 15(2):245–295.

Ashcraft, A., Fernández-Val, I., and Lang, K. (2013). The consequences of teenage childbearing: Consistent estimates when abortion makes miscarriage nonrandom. *Economic Journal*, 123(571):875–905.

Baird, S., McIntosh, C., and Özler, B. (2011). Cash or condition? evidence from a cash transfer experiment. *Quarterly Journal of Economics*, 126(4):1709–1753.

Black, S. E., Devereux, P. J., and Salvanes, K. G. (2008). Staying in the classroom and out of the maternity ward? the effect of compulsory schooling laws on teenage births. *Economic Journal*, 118(530):1025–1054.

Bleakley, H. (2010). Malaria eradication in the Americas: A retrospective analysis of childhood exposure. *American Economic Journal: Applied Economics*, 2(2):1–45.

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. H. C. (2024). Difference-in-differences with a continuous treatment. NBER Working Paper 32117, National Bureau of Economic Research.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.

Cutler, D., Fung, W., Kremer, M., Singhal, M., and Vogl, T. (2010). Early-life malaria exposure and adult outcomes: Evidence from malaria eradication in India. *American Economic Journal: Applied Economics*, 2(2):72–94.

Duflo, E. (2001). Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment. *American Economic Review*, 91(4):795–813.

Duflo, E., Dupas, P., and Kremer, M. (2015). Education, HIV, and early fertility: Experimental evidence from Kenya. *American Economic Review*, 105(9):2757–2797.

Field, E. and Ambrus, A. (2008). Early marriage, age of menarche, and female schooling attainment in Bangladesh. *Journal of Political Economy*, 116(5):881–930.

Fletcher, J. M. and Wolfe, B. L. (2009). Education and labor market consequences of teenage childbearing: Evidence using the timing of pregnancy outcomes and community fixed effects. *Journal of Human Resources*, 44(2):303–325.

Frakes, M. D. and Harding, M. C. (2015). The effect of statutory rape laws on teen birth rates. *American Law and Economics Review*, 17(2):409–461.

Glynn, J. R., Kayuni, N., Floyd, S., Banda, E., Francis-Chizororo, M., Tanton, C., Molesworth, A., Hemmings, J., Crampin, A. C., and French, N. (2010). Age at menarche, schooling, and sexual debut in northern Malawi. *PLoS ONE*, 5(12):e15334.

Glynn, J. R., Sunny, B. S., DeStavola, B., Dube, A., Chihana, M., Price, A. J., and Crampin, A. C. (2018). Early school failure predicts teenage pregnancy and marriage: A large population-based cohort study in northern Malawi. *PLoS ONE*, 13(5):e0196041.

Government of Malawi (2011). Penal Code (Amendment) Act, 2011 and consolidated Penal Code provisions on defilement and sexual activity with a child. Act No. 1 of 2011; consolidated text via MalawiLII.

Government of Malawi (2015). Marriage, Divorce and Family Relations Act. Act No. 4 of 2015; consolidated text via MalawiLII.

Government of Malawi (2023). Penal Code (Amendment) Act, 2023. Act No. 8 of 2023; MalawiLII.

Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.

Jayachandran, S. (2015). The roots of gender inequality in developing countries. *Annual Review of Economics*, 7(1):63–88.

Kangaude, G. D. (2017). Adolescent sex and "defilement" in Malawi law and society. *African Human Rights Law Journal*, 17(2):527–549.

Kangaude, G. D. (2021). Positive approaches to childhood sexuality and transforming gender norms in Malawi. In Durojaye, E., Mirugi-Mukundi, G., and Ngwena, C., editors, *Advancing Sexual and Reproductive Health and Rights in Africa: Constraints and Opportunities*, pages 99–114. Routledge.

Kangaude, G. D. and Skelton, A. (2018). (de)criminalizing adolescent sex: A rights-based assessment of age of consent laws in eastern and southern Africa. *SAGE Open*, 8(4):1–12.

Keane, M. P. and Neal, T. (2024). A practical guide to weak instruments. *Annual Review of Economics*, 16(1):185–212.

Lepage, L.-P. (2022). Do age of consent laws decrease teen births? *Canadian Journal of Economics/Revue canadienne d'économique*, 55(3):1431–1459.

Levine, P. B. (2014). Teenage childbearing and labor market implications for women. IZA World of Labor, article 28.

Lucas, A. M. (2010). Malaria eradication and educational attainment: Evidence from Paraguay and Sri Lanka. *American Economic Journal: Applied Economics*, 2(2):46–71.

Mussa, R. (2016). Labour market transitions of young women and men in Malawi: Results of the 2014 school-to-work transition survey. Work4Youth Publication Series 35, International Labour Office, Geneva.

National Statistical Office (NSO) [Malawi] and ICF (2017). Malawi Demographic and Health Survey 2015–16. Technical report, National Statistical Office and ICF, Zomba, Malawi and Rockville, Maryland, USA.

National Statistical Office (NSO) [Malawi] and ICF Macro (2011). Malawi Demographic and Health Survey 2010. Technical report, National Statistical Office and ICF Macro, Zomba, Malawi and Calverton, Maryland, USA.

National Statistical Office (NSO) [Malawi] and ORC Macro (2001). Malawi Demographic and Health Survey 2000. Technical report, National Statistical Office and ORC Macro, Zomba, Malawi and Calverton, Maryland, USA.

National Statistical Office (NSO) [Malawi] and ORC Macro (2005). Malawi Demographic and Health Survey 2004. Technical report, National Statistical Office and ORC Macro, Calverton, Maryland, USA.

National Statistical Office (NSO) [Malawi], Utica International, and ICF (2026). Malawi Demographic and Health Survey 2024: Final Report. Technical report, National Statistical Office, Utica International, and ICF, Zomba, Malawi; Columbia, Maryland, USA; and Rockville, Maryland, USA.

Nketiah, A. (2026). Raising the bar: Age of consent reform and women's outcomes in Malawi. Unpublished manuscript.

Psaki, S. R., Chuang, E. K., Melnikas, A. J., Wilson, D. B., and Mensch, B. S. (2019). Causal effects of education on sexual and reproductive health in low and middle-income countries: A systematic review and meta-analysis. *SSM - Population Health*, 8:100386.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in Stata using boottest. *Stata Journal*, 19(1):4–60.

Soler-Hampejsek, E., Kangwana, B., Austrian, K., Amin, S., and Psaki, S. R. (2021). Education, child marriage, and work outcomes among young people in rural Malawi. *Journal of Adolescent Health*, 69(6S):S57–S64.

Sunny, B. S., DeStavola, B., Dube, A., Price, A., Kaonga, A. M., Kondowe, S., Crampin, A. C., and Glynn, J. R. (2019). Lusting, learning and lasting in school: Sexual debut, school performance and dropout among adolescents in primary schools in karonga district, northern Malawi. *Journal of Biosocial Science*, 51(5):720–736.

Sunstein, C. R. (1996). On the expressive function of law. *University of Pennsylvania Law Review*, 144(5):2021–2053.

UN Women (2024). Gender pay gap and labour market inequalities in Malawi. Technical report, UN Women East and Southern Africa Regional Office, Nairobi, Kenya.

UNICEF Malawi (2019). Access to criminal justice services: The case of survivors of sexual violence in Malawi. Technical report, UNICEF Malawi, Lilongwe, Malawi.

Wado, Y. D., Sully, E. A., and Mumah, J. N. (2019). Pregnancy and early motherhood among adolescents in five east African countries: A multi-level analysis of risk and protective factors. *BMC Pregnancy and Childbirth*, 19(1):59.

Wilson, N. (2022). Child marriage bans and female schooling and labor market outcomes: Evidence from natural experiments in 17 low- and middle-income countries. *American Economic Journal: Economic Policy*, 14(3):449–477.

World Bank (2021). Unlocking Malawi's economic growth by bridging the widening gender gaps in the labor workforce. Technical report, World Bank, Washington, DC.

A Appendix Tables

Table A1: Treatment and Outcome Measurement

Measure	Unit	Use in analysis
First sex before age 16	0/1	Endogenous treatment
Skilled occupation	0/1	Primary occupational outcome
Occupational status index	0–5 index	Primary occupational outcome
Professional-or-clerical occupation	0/1	Secondary occupational outcome
Currently working	0/1	Employment-selection outcome

Notes: Binary variables are in probability units. The occupational status index ranges from 0 to 5.

Table A2: Cell-Level Historical Intensity

Measure	Cells	Mean	SD	Min	Max
Historical early-debut share, born $\leq$ 1970 (%)	72	38.778	11.796	10.981	66.480

Notes: Summary statistics are reported at the district-by-ethnicity cell level. Cells have at least 25 valid observations. Intensity is measured using cohorts born in 1970 or earlier. DHS sampling weights are used.

Table A3: Historical Intensity by Birth Year

Birth year	Cohort group	N	Mean intensity (%)	Mean std. intensity
1987	Control	2,137	38.08	-0.059
1988	Control	2,372	37.98	-0.068
1989	Control	1,779	38.14	-0.054
1990	Control	1,957	38.87	0.008
1991	Control	1,797	37.62	-0.098
1992	Control	2,318	37.93	-0.072
1993	Control	2,202	37.71	-0.091
1994	Control	2,325	38.06	-0.061
1999	Exposed	1,469	38.78	0.000
2000	Exposed	942	38.85	0.006
2001	Exposed	591	39.47	0.059
2002	Exposed	632	39.02	0.021
2003	Exposed	821	38.80	0.002
2004	Exposed	902	38.97	0.016
2005	Exposed	825	40.12	0.114
2006	Exposed	770	38.98	0.017
Control cohorts	1987–1994	16,887	38.04	-0.063
Exposed cohorts	1999–2006	6,952	39.08	0.025

Notes: DHS sampling weights are used. Donut cohorts are excluded. Historical intensity is assigned at the district-by-ethnicity level using cohorts born in 1970 or earlier. Test of equality across birth-year indicators: $F(15, 23823) = 4.51$, $p < 0.001$.

Table A4: Ethnicity Composition and Historical Intensity

Ethnicity	Sample share (%)	Mean intensity (%)	District-by-ethnicity cells	Districts
Chewa	37.2	32.48	18	18
Lomwe	17.8	49.64	10	10
Yao	14.0	49.90	10	10
Ngoni	10.7	32.41	11	11
Tumbuka	9.8	28.88	8	8
Sena	3.5	32.74	3	3
Mang'anja	2.7	47.43	8	8
Tonga	1.3	39.98	2	2
Other	1.1	40.70	10	10
Nkhonde	1.0	38.92	6	6

Notes: DHS sampling weights are used. Donut cohorts are excluded. Historical intensity is assigned at the district-by-ethnicity level using cohorts born in 1970 or earlier. The main analysis contains 72 district-by-ethnicity cells and 27 districts.

Table A5: Employment Selection and Occupational Outcomes Among Workers

Panel A. Reduced-form estimate for current employment			
	Currently working		
Exposed cohort x standardized intensity	0.0135		
	(0.0260)		
Observations	20,047		
Employment rate	0.643		
Panel B. Reduced-form occupational estimates among currently working women			
	Skilled	Status index	Professional-or-clerical
Exposed cohort x standardized intensity	0.0415**	0.162**	0.0414***
	(0.0199)	(0.0704)	(0.0153)
Observations	12,233	12,233	12,233
Outcome mean	0.167	1.670	0.099
Panel C. 2SLS estimates among currently working women			
	Skilled	Status index	Professional-or-clerical
Early sexual debut (instrumented)	-0.775*	-3.026**	-0.773**
	(0.416)	(1.494)	(0.341)
Observations	12,233	12,233	12,233
Outcome mean	0.167	1.670	0.099
KP weak-ID F	24.18	24.18	24.18

Notes: Standard errors are in parentheses. Panel A uses current employment as the dependent variable. Panel B reports reduced-form coefficients on the reform-intensity interaction. Panel C reports 2SLS coefficients on early sexual debut instrumented by the reform-intensity interaction. District-by-ethnicity, birth-year, and survey-year fixed effects are included. DHS sampling weights are used. Standard errors are clustered by district-by-ethnicity cell. KP weak-ID F is the Kleibergen-Paap rk Wald weak-identification F statistic. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: Reduced-Form and 2SLS Robustness Specifications

Panel A. Reduced-form robustness

Specification	Skilled	Status index	Professional-or-clerical	N
Add household wealth control	0.0295***	0.134***	0.0249***	19,307
	(0.00904)	(0.0370)	(0.00704)	
Binary intensity instrument	0.0659***	0.294***	0.0471***	19,307
	(0.0170)	(0.0783)	(0.0140)	
Minimum age 22	0.0380***	0.168***	0.0322***	12,336
	(0.0134)	(0.0433)	(0.0102)	
Exclude major urban centers	0.0219***	0.0986***	0.0186***	17,745
	(0.00827)	(0.0306)	(0.00673)	

Panel B. 2SLS robustness

Specification	Skilled	Status index	Professional-or-clerical	KP F	N
Add household wealth control	-0.674***	-3.065***	-0.570***	22.33	19,307
	(0.195)	(0.787)	(0.169)		
Binary intensity instrument	-0.676***	-3.010***	-0.483***	24.00	19,307
	(0.175)	(0.726)	(0.148)		
Minimum age 22	-0.572***	-2.521***	-0.484***	27.89	12,336
	(0.156)	(0.595)	(0.140)		
Exclude major urban centers	-0.444***	-1.998***	-0.377***	23.05	17,745
	(0.150)	(0.453)	(0.132)		

Notes: Standard errors are in parentheses. Panel A reports reduced-form coefficients on the reform-intensity interaction. Panel B reports 2SLS coefficients on early sexual debut instrumented by the same interaction. District-by-ethnicity, birth-year, and survey-year fixed effects are included. DHS sampling weights are used. Standard errors are clustered by district-by-ethnicity cell. Binary outcomes are in probability units; occupational status is a 0–5 index. KP F is the Kleibergen-Paap rk Wald weak-identification F statistic. The binary-intensity row uses exposed cohort interacted with above-median historical intensity; other rows use standardized historical intensity. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: Construction of Occupational Outcome Measures

DHS code	DHS occupation category	Prof./ clerical	Skilled	Status	Status-index category
0	Not currently working	No	No	0	Not currently working
1	Professional, technical, or managerial	Yes	Yes	5	Professional, technical, or managerial
2	Clerical	Yes	Yes	4	Clerical
3	Sales	No	No	2	Sales/domestic/lower-service work
4	Agricultural, self-employed	No	No	1	Agriculture or unskilled manual work
5	Agricultural, employee	No	No	1	Agriculture or unskilled manual work
6	Household and domestic work	No	No	2	Sales/domestic/lower-service work
7	Services	No	Yes	3	Services or skilled manual work
8	Skilled manual work	No	Yes	3	Services or skilled manual work
9	Unskilled manual work	No	No	1	Agriculture or unskilled manual work
10	Armed forces	–	–	–	Not assigned
96	Other occupation	–	–	–	Not assigned

Notes: Professional-or-clerical equals one for clerical, professional, technical, and managerial work. Skilled occupation equals one for services, skilled manual, clerical, professional, technical, and managerial work. The occupational status index ranges from 0 to 5. Armed forces and other occupation codes are excluded.

Table A8: Sample Construction

Restriction	Control	Exposed	Total
Main cohorts, age ≥ 16, first-stage data available	16,887	6,952	23,839
Restrict to age ≥ 18	14,235	5,812	20,047
Valid occupational classification	13,868	5,439	19,307

Notes: Counts are unweighted. The final row is the main occupational estimation sample.

B Appendix Figures

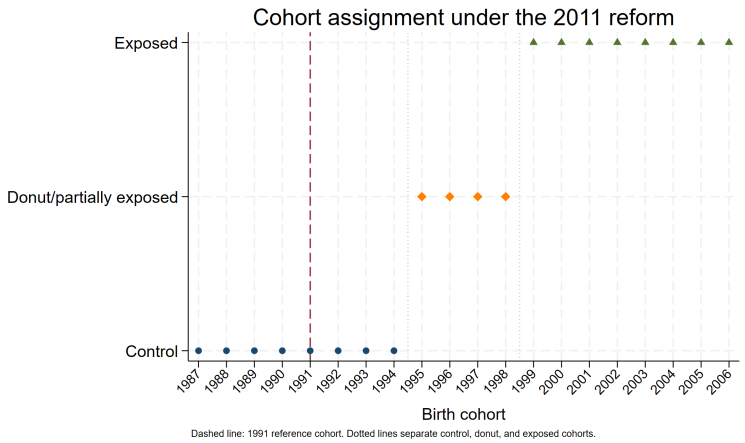

Figure A1: Cohort assignment under the 2011 reform.

Notes: Control cohorts are born in 1987–1994. Exposed cohorts are born in 1999–2006. Donut cohorts are born in 1995–1998.

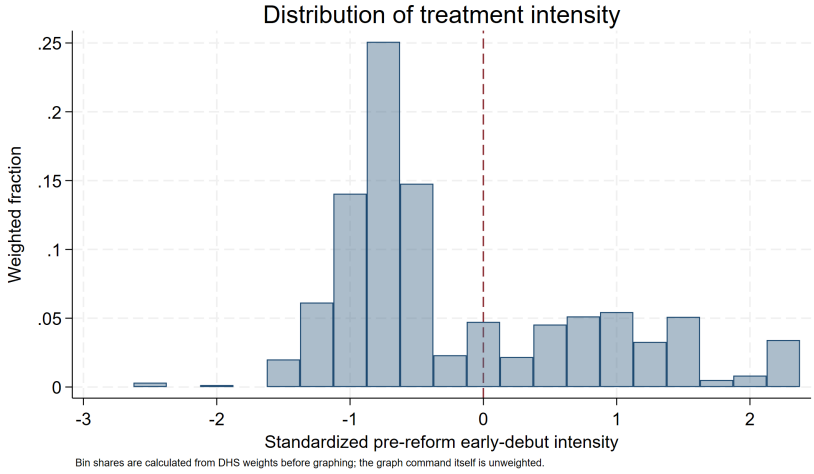

Figure A2: Distribution of standardized historical early-debut intensity.

Notes: Historical intensity is measured at the district-by-ethnicity level using cohorts born in 1970 or earlier. The standardized measure has mean zero and unit standard deviation. The main analysis uses 72 district-by-ethnicity cells.

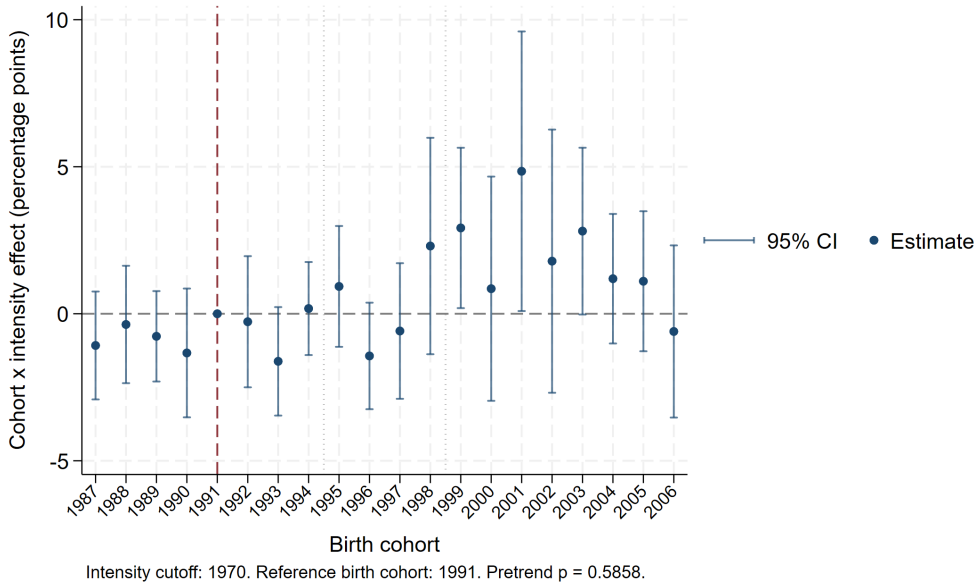

Figure A3: Cohort-by-intensity diagnostic profile for the professional-or-clerical occupation outcome.

Notes: Coefficients are birth-cohort interactions with standardized historical early-debut intensity relative to the omitted 1991 birth cohort. Coefficients are in percentage points. Vertical bars report 95% confidence intervals. Joint pretrend test $p = 0.5858$.

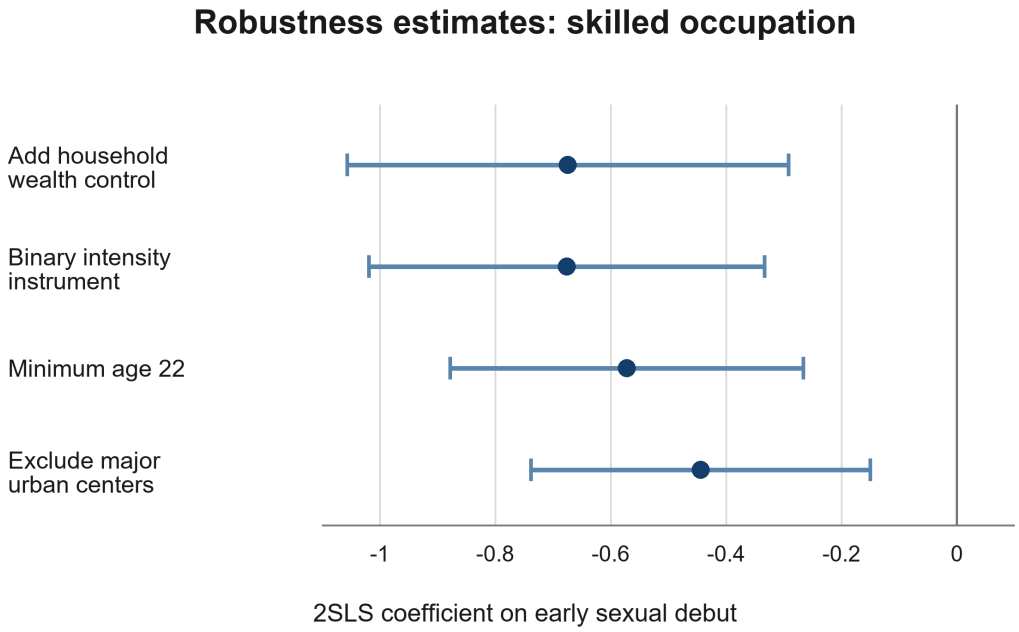

Figure A4: Robustness estimates for skilled occupation.

Notes: Estimates are 2SLS coefficients for early sexual debut. Horizontal bars report 95% confidence intervals. Specifications correspond to the 2SLS columns in Appendix Table A6.

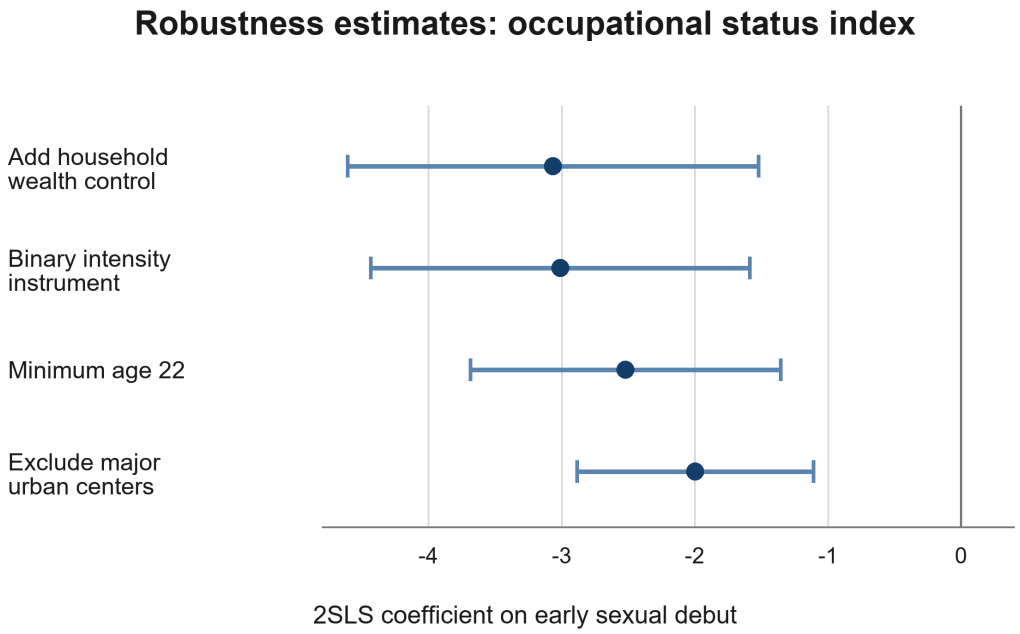

Figure A5: Robustness estimates for the occupational status index.

Notes: Estimates are 2SLS coefficients for early sexual debut. The outcome is measured on the 0–5 occupational-status scale. Horizontal bars report 95% confidence intervals. Specifications correspond to the 2SLS columns in Appendix Table A6.

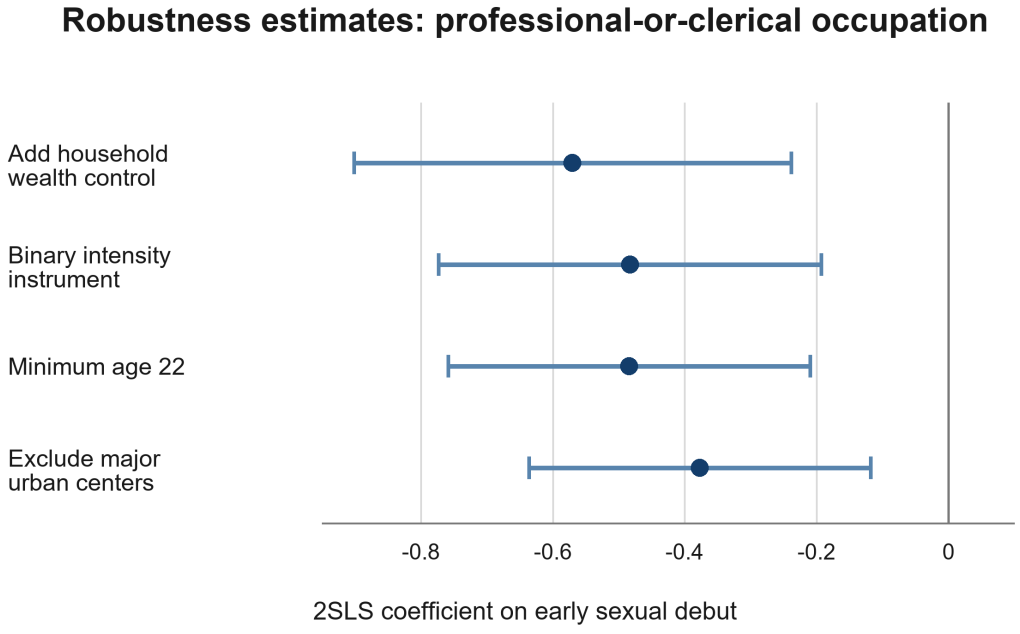

Figure A6: Robustness estimates for the professional-or-clerical occupation outcome.

Notes: Estimates are 2SLS coefficients for early sexual debut. Coefficients are in probability units. Horizontal bars report 95% confidence intervals. Specifications correspond to the 2SLS columns in Appendix Table A6.