HOUSE OF SUSTAINABLE SOCIETY

HOSS
WORKING PAPER SERIES

NO. 2026-01, APRIL 2026

IDENTITY AND HUMAN CAPITAL INVESTMENT: EVIDENCE FROM VEILING BAN REMOVAL IN TURKEY

Merve Demirel
Institute for International Economic Studies,
Stockholm University

Avenia Ghazarian
House of Sustainable Society,
Stockholm School of Economics

Identity and Human Capital Investment: Evidence from Veiling Ban Removal in Turkey*

Merve Demirel† Avenia Ghazarian‡

April, 2026

Abstract

This paper examines how restrictions on religious expression affect women's educational attainment. We study the 2010 removal of the headscarf ban in Turkish universities, which had long limited access to higher education for visibly religious women. Our empirical strategy combines cohort-level variation in exposure to the reform with individual-level variation in the propensity to veil within a difference-in-differences framework. We estimate veiling propensities using an early wave of the Turkish Demographic and Health Survey and predict them for a later sample using both machine learning and parametric methods. We show that lifting the ban significantly increased educational attainment among women with a higher propensity to veil. These gains appear to be concentrated around the transition into and progression through secondary school. The results remain similar when, instead of individual-level propensities, we use pre-reform veiling prevalence at the province level as an alternative exposure measure.

Keywords: Identity; Religious expression; Veiling ban; Turkey

*We are grateful to Martina Björkman Nyqvist, Johan Callermo, Tore Ellingsen, Akib Khan, Sina Sheikholeslami, Abhijeet Singh, Anna Tompsett, and seminar participants at Stockholm School of Economics and the IIES for valuable comments and helpful discussions.

†Stockholm University and the IIES, merve.demirel@iies.su.se

‡Stockholm School of Economics, avenia.ghazarian@hhs.se

1 Introduction

Policymakers use a variety of tools in order to define and regulate the role of religion in public life. A very prominent example is the regulation of Islamic veil: across countries, policymakers have sought to regulate veiling, either by restricting it or by making it compulsory.[1] These regulations, regardless of direction, impose practical and symbolic constraints on women's access to public spaces and economic participation. By requiring or restricting specific forms of dress and limiting how individuals express identity and signal group membership, they unavoidably link economic decisions to identity expression.

Educational institutions have been central sites for the enforcement of veiling regulations. Theoretically, to the extent that veiling serves as an expression of identity, regulating it through the educational system may create a trade-off between identity expression and educational attainment: those who do not wish to comply with policies restricting participation may choose to drop out of school or avoid educational institutions altogether (Carvalho et al., 2024). At the same time, if schools are effective at transmitting the intended norms, whether secular or religious, exposure to these environments may reduce the salience of veiling as an identity marker and limit resistance to such policies.[2] In addition, veiling may also reflect compliance with household or community norms rather than individual preference (Akerlof and Kranton, 2002; Bursztyn, González, et al., 2020; Maurin and Navarrete H, 2023), in which case, restrictions on veiling could reduce the cost of attending school for some women by weakening externally imposed pressures. The net effect of veiling bans on educational attainment is therefore theoretically ambiguous, and ultimately an empirical question.

We investigate this question in Turkey, where bans on veiling in educational institutions and public sector employment were a defining feature of the secular state for decades and effectively restricted veiled women from accessing tertiary education and public sector employment. Throughout the ban period, an estimated 5,000 women wearing headscarves were dismissed from public institutions for violating the dress code, around 10,000 were pressured to resign, and roughly 3,000 female students were expelled from universities for noncompliance with the ban (Albayrak, 2023). Beginning in the early 2010s, these restrictions were gradually lifted, allowing veiled women to attend university, work in public institutions, and study in secondary schools without removing their headscarves.

[1] For example, veiling became mandatory for women in Iran after the 1979 Islamic Revolution, while France banned headscarves in public schools in 2004. Belgium and the Netherlands later introduced bans on face covering in public spaces (Carvalho, 2013).

[2] Bisin and Verdier, 2000; Bisin and Verdier, 2001 provide a useful theoretical framework for this channel.

In this context, we ask whether lifting restrictions on religious expression affects women's educational investments.

Understanding the relationship between religious expression and educational investments is empirically challenging. Religious expression is typically endogenous to schooling choices: individuals who place greater value on religious expression may sort into educational environments that accommodate it. At the same time, measures of religious expression are rarely observed in standard datasets, which makes these mechanisms difficult to study empirically. Moreover, veiling often displays little variation within a context: some settings, such as Egypt or Iran, are characterized by near-universal adoption, while in others prevalence is very low, leaving little identifying variation. As a result, existing evidence on the effects of veiling regulations on human capital is limited and largely drawn from contexts in which veiled women constitute a relatively small minority.

Turkey offers a compelling setting in which to examine this issue. The removal of the veiling ban in universities represents a clear and well-defined institutional change in the conditions under which veiled women could access and remain in education. At the same time, the social relevance of the policy is potentially large, given the widespread prevalence of veiling: 66% of women reported wearing a headscarf in 2010, at the time of ban removal. Finally, the availability of nationally representative survey data in Turkey that directly record women's veiling behavior helps overcome a key data limitation in the existing literature.

The removal of the veiling ban in universities has likely been most consequential for women who would choose to veil, as it is precisely this group whose schooling decisions might have been constrained by the ban. Our identification strategy therefore leverages two sources of variation in a difference-in-differences framework. First, the timing of the reform which generates variation across cohorts in exposure to the policy change. Second, individual-level variation in one's likelihood of being veiled.

We primarily rely on the two most recent waves of the Demographic and Health Surveys (DHS) in Turkey conducted in 2013 and 2018. A key feature of the 2013 wave for our purposes is that it contains individual-level information on whether a woman veils. This information, however, is not available in the 2018 wave, from which we construct our main analysis sample. To address this limitation, we use the 2013 wave to estimate a model of veiling and predict veiling probabilities for women in the 2018 sample. Our main prediction model is a random forest, and we further validate the resulting estimated probabilities using parametric approaches.

Overall, we find that the removal of the ban increased educational attainment among women who were more likely to veil. Women in post-reform cohorts gained between 0.6 and 0.9 additional years of schooling relative to older cohorts unaffected by the reform. This corresponds to a 6 to 10 percent increase over a baseline mean of 8.9 years. The event-study analysis supports this interpretation. The estimated coefficients are close to zero for older, unaffected cohorts and turn positive and grow in magnitude for cohorts exposed to the policy change.

The findings further suggest that the gains are concentrated around the transition into and within secondary schooling, suggesting that the ban had been most binding at this stage of the educational ladder. Corroborating these findings, we obtain similar results when exposure to the reform is measured using pre-reform veiling rates at the province level rather than individual predicted probabilities.

A natural concern with our approach is that veiling probabilities in the main analysis sample are estimated rather than directly observed. This raises two related issues: first, whether the results are sensitive to how predicted probabilities are converted into binary veiling status, and second, whether treating predicted status as known with certainty overstates precision. We address the first by showing that the positive effects hold across a wide range of threshold cutoffs. For the second, we implement a simulation-based multiple imputation procedure in the spirit of Rubin, 1987, that assigns veiling status probabilistically, incorporating classification uncertainty into the estimation. The results remain positive under this more conservative approach, though somewhat less precisely estimated. Importantly, we show that our findings do not hinge on the predicted probabilities at all. In an alternative specification, where we set aside individual-level veiling predictions entirely and instead use pre-reform veiling rates at the province level as a measure of exposure, we find that the schooling gains are concentrated among women from provinces in the top quintile of pre-reform veiling prevalence.

This paper relates to the literature examining how identity concerns shape economic behavior (Akerlof and Kranton, 2000; Akerlof and Kranton, 2002). A central insight of this literature is that economic choices are not made solely on the basis of material returns, but also in relation to the social meanings and identity content attached to them. In the labor market, identity can either constrain or enable economic behavior depending on the social context. For example, Oh (2023) shows that workers in rural India forgo substantial payments to avoid job offers associated with other castes, while Shofia (2022) finds that Indonesian women adopt the veil as a signaling device that facilitates labor market entry

while remaining consistent with identity-related norms. In education, Bursztyn, Egorov, et al. (2019) show that social identity shapes educational effort: students reduce effort when high achievement conflicts with identity-congruent behavior, reflecting concerns about peer perception. Delavande and Zafar (2019) further document that ideological alignment is a first-order determinant of university choice in Pakistan, outweighing expected labor market returns. [3]

Second, this paper provides empirical support for the theoretical framework proposed by Carvalho et al. (2024), which studies how top-down modernization policies can generate cultural backlash. Consistent with Abdelgadir and Fouka (2020), who show that France's veiling ban decreased Muslim girls' secondary schooling, our findings speak directly to the mechanism at the core of Carvalho et al. (2024): when state-led reforms are perceived as threatening group identity, they may induce behavioral responses that run counter to the policy's intended goals.[4]

Third, this paper contributes to a growing literature on the relationship between religiosity and human capital investment, where existing causal evidence points to a tradeoff between the two. Gulesci and Meyersson (2015) show that an additional year of schooling in Turkey reduces religiosity, while Kökkizil (2022) finds that exposure to Ramadan at the time of primary school enrollment lowers girls' enrollment. In this setting, institutional arrangements play a central role in shaping girls' schooling decisions. Meyersson (2014) finds that the election of Islamic mayors in 1994 improved schooling outcomes for daughters from poor and pious families, driven by an expansion of educational facilities sponsored by religious charities that may have lowered the cultural cost of schooling for conservative households. A more closely related set of papers examines the effects of veiling regulations on women's labor market outcomes in Turkey. Corekcioglu (2021) and Yurdakul and Lu (2025) document increases in employment among veiled women following the removal of the public-sector veiling ban in 2013. By contrast, Uğur (2020) studies the effect of the veiling ban on university access using a regression discontinuity design and finds limited impacts.[5]

[3] A related perspective is offered by the aspirations literature, which emphasizes that goals are socially formed and shaped by one's reference group. In this framework, identity can influence educational investment by shaping the set of futures individuals perceive as attainable or desirable (Appadurai, 2004; Dalton et al., 2016; Genicot and Ray, 2017).

[4] This backlash mechanism also connects to a broader literature on norm change through institutional exposure, including work showing that school curricula can reshape gender norms (e.g., Dhar et al., 2022). Together, these studies suggest that the identity content and framing of education, not only access to it, can shape downstream behavioral outcomes.

[5] Her analysis focuses on cohorts born between 1973 and 1987, when overall university attainment remained

The remainder of the paper is structured as follows. Section 2 provides background on the origins and subsequent removal of the veiling ban in Turkey and presents descriptive evidence on veiling and educational attainment. Section 3 describes the data. Section 4 outlines the empirical strategy for predicting veiling and identifying the effects of lifting the ban on women's educational outcomes. Section 5 presents and interprets the results. Finally, Section 6 concludes with a discussion.

2 Background

2.1 The removal of the veiling ban

Since the establishment of the Republic of Turkey in 1923 by Mustafa Kemal Atatürk, a series of reforms aimed at secularizing the nation were implemented.[6] Among these, several reforms targeted the role of religious symbols and dress in public life. The Law on the Prohibition of the Wearing of Certain Garments (1934) outlawed religious attire — including the *sarık* and *cübbe*, traditionally worn by religious clergy — except during worship.[7] The adoption of the *Medeni Kanun* (Civil Code, 1926), modeled on the Swiss Civil Code, further secularized family and personal law, freeing women from religiously prescribed dress obligations. While the headscarf was never formally banned during this period, it was widely portrayed as a symbol of backwardness and as incompatible with the modern national identity, and women were actively encouraged to adopt Western-style clothing (Albayrak, 2023; Zurcher, 2004). Debates over religion in public life, and the headscarf in particular, continued for decades, with formal nationwide restrictions on religious dress in public institutions emerging only after the 1980 military coup.

Following the 1980 military coup, Turkey's 1982 constitution established regulations governing the dress code of civil servants, prohibiting the display of religious and political symbols in public offices, which included requirements for women to uncover their heads. [8] These restrictions were justified as necessary to uphold the principle of secularism

very low, at roughly 3.5 percent. By the time the ban was lifted in 2010, educational attainment had expanded substantially.

[6] Some of the reforms include: adoption of European constitutional and legal standards (Balci and Monceau, 2023), the closure of dervish lodges and zawiyas, tombs and madrasas that provided religious education (Albayrak, 2023), and the Turkification of the azan (Albayrak, 2023).

[7] See https://www.mevzuat.gov.tr/MevzuatMetin/1.3.2596.pdf.

[8] See *Kamu Kurum ve Kuruluşlarında Çalışan Personelin Kılık ve Kıyafetine Dair Yönetmelik* Council of Ministers Decision No. 8/5105, *Resmî Gazete* [Official Gazette], Oct. 25, 1982, No. 17849. https://www.resmigazete.gov.tr/arsiv/17849.pdf

(laiklik) by removing religious symbols from state institutions and to ensure equality across religious affiliations and genders. In 1997, following a National Security Council decision calling for strict, uncompromising enforcement of dress code laws across public institutions, YÖK (Council of Higher Education) issued a directive requiring universities to deny entry to headscarf-wearing students, extending and hardening enforcement of dress code restrictions at the university level.

Throughout the years, several pieces of legislation aimed at removing the headscarf ban in universities and government offices were passed, but they were annulled by the Constitutional Court on the grounds that they violated the principle of secularism. The political landscape shifted after the Justice and Development Party (AKP), founded in 2001, came to power in the 2002 general elections with 34.3% of the vote, forming a single-party government with 363 of the 550 seats in parliament. After assuming power, the AKP initiated policies that gradually increased the visibility of religion in the public sphere. In 2010, the party introduced a significant package of constitutional amendments that limited the ability of the Constitutional Court to oppose regulatory changes, which gained broad public approval through a referendum. Buoyed by this public support, the ban on headscarves was first lifted in universities (2010), then in government offices (2013), and in high schools (2014).

In practice, the enforcement of the ban was delegated to universities; allowing university presidents, deans, and faculty to determine the extent of its application. As a result, implementation varied across institutions and over time: while some universities enforced the ban strictly, others applied it more leniently (Cindoğlu, 2011). Qualitative accounts nevertheless suggest that students wearing headscarves faced sustained pressure throughout their education. In interviews with veiled university graduates, many described repeatedly negotiating these restrictions in order to attend classes or access university facilities (Cindoğlu, 2011).

Veiled students adopted a range of strategies to circumvent these restrictions, including wearing wigs, temporarily removing their headscarves, or transferring majors or institutions; whether such strategies were tolerated depended largely on the university. This variation created significant uncertainty and arbitrariness for veiled women (Cindoğlu, 2011). As a result, the issue remained a central consideration for those aspiring to pursue post-secondary education, and access to higher education could be effectively constrained when such coping strategies were unavailable or socially unacceptable. This constraint was potentially substantial given that a majority of women in Turkey wear some form of head

covering (Çarkoğlu and Toprak, 2007; KONDA, 2019). Consistent with this concern, survey data from 2006 show that 24.5% of parents would not approve of their daughters removing their headscarves to attend university (Çarkoğlu and Toprak, 2007), suggesting that the ban may have restricted educational access for a significant portion of the population.

2.2 Education and veiling in Turkey

Veiling is widespread in Turkey: approximately 65% of women use some form of head or body covering in daily life (Figure 1), with headscarves being the most prevalent at around 55%. The headscarf, which covers only the head while leaving some hair visible, is distinct from the *çarşaf* or chador, which covers the entire body from head to toe except for the eyes, and from the turban, which covers the full head, neck, and shoulders with no hair visible but leaves the face bare. [9] (Çarkoğlu and Toprak, 2007) The practice is also geographically diffuse: while coverage rates vary across provinces, veiling remains common even in the least observant regions, where roughly 30% of women still veil (see Figure 5). While older women are slightly more likely to veil, this pattern likely reflects cohort differences rather than women adopting the veil as they age. Few respondents cite age as a primary reason for veiling (Çarkoğlu and Toprak, 2007), and veiling rates are broadly stable across the age distribution (Figure B3).

Veiling is negatively associated with educational attainment. On average, veiled women have 5.9 years of schooling, which is 4.3 fewer years than unveiled women. Figure B1 plots years of schooling by age for veiled and unveiled women separately, and shows that this pattern persists across the age distribution. [10] Veiled women are also more likely to come from rural areas and lower socioeconomic backgrounds, and report holding more harmful gender norms (Table 1), all of which independently predict lower educational attainment.

However, these average differences mask substantial heterogeneity among veiled women. Veiling remains common even among highly educated women. Figure B1 plots the share of veiled women by educational attainment and shows that a non-negligible fraction of women with tertiary education continue to veil. Attitudes toward gender roles also differ on average between veiled and unveiled women, but the relationship is not uniform. Figure B4 plots the composition of veiled and unveiled women across quintiles of acceptance of wife-beating. The share of veiled women rises with acceptance of wife-beating, exceeding

[9] Definitions follow the survey classifications in Çarkoğlu and Toprak (2007). Crucially, the headscarf ban applied to all forms of veiling.

[10] Note that veiling was banned in primary and secondary schools throughout the study period for women in our sample.

80% in the top quintile. However, in the bottom quintile (where wife-beating is least accepted), veiled and unveiled women are roughly equally represented, indicating that a substantial share of veiled women do not hold more regressive gender attitudes than their unveiled counterparts.

Overall, these patterns suggest considerable heterogeneity among veiled women. Although veiling is correlated with characteristics associated with lower educational attainment, it is also common among women who do not exhibit strongly traditional gender attitudes and who attain high levels of education. This heterogeneity implies the existence of a subgroup of veiled women with high educational aspirations. Restrictions on veiling in higher education are most likely to affect this group, for whom educational attainment would otherwise be feasible.

3 Data

We use data from the two most recent rounds of the Demographic and Health Survey (DHS) for Turkey, conducted in 2013 and 2018. The DHS surveys are nationally representative and provide rich information on demographic, socioeconomic, and household characteristics. Our analysis focuses on the women's module, which surveys women aged 15 to 49. Importantly, the 2013 wave includes a question on veiling behavior, allowing us to directly observe whether a respondent reports wearing a headscarf. [11] In contrast, the 2018 wave does not include this question. Table 1 presents descriptive statistics for veiled (column 1) and unveiled (column 2) women in the 2013 survey wave across the three dimensions of demographic characteristics, educational attainment, and attitudes toward wife beating. [12]

4 Empirical Strategy

4.1 Veil prediction

To investigate whether the removal of the veiling ban affects schooling outcomes of women, we exploit two sources of variation. First, we use individual-level differences in veiling

[11] Earlier DHS waves for Turkey, including 2003 and 2008, also contain a question on veiling. However, these waves survey ever-married women only. To avoid potential selection concerns arising from changes in marital composition across cohorts and over time, we restrict the construction of our veiling-based measure to the 2013 wave, which covers all women aged 15–49 regardless of marital status.

[12] In the 2013 DHS, 106 out of 9,630 respondents (1.1%) report wearing a headscarf "irregularly". We consider this group as veiled throughout the analysis.

behavior, defined as the likelihood of a woman to wear a headscarf when outside the home. Second, we exploit variation in exposure to the reform across birth cohorts.

As noted above, the 2018 DHS wave, which we use to measure educational outcomes for the cohorts of interest, does not include information on women's veiling behavior. Hence, we cannot directly classify respondents in 2018 as veiled or unveiled. To address this limitation, we exploit the 2013 DHS wave, which does contain a question on headscarf use, to train a supervised classification model. Specifically, we estimate a random forest model that learns the relationship between observable characteristics and reported veiling status in 2013. We then apply the trained model to the 2018 wave to generate predicted probabilities of veiling for each respondent. Importantly, this approach relies on the underlying assumption that the relationship between observable characteristics and veiling behavior remains sufficiently stable over time, such that a model trained on the 2013 cross-section can meaningfully predict veiling status in 2018. In particular, the strategy requires that the relative ordering implied by the model be preserved: women predicted to have a higher propensity to veil in 2013 should remain relatively more likely to veil in 2018.

To assess the plausibility of this assumption, we rely on evidence from the reports on KONDA Barometer surveys, which are nationally representative opinion surveys conducted in Turkey and regularly report on women's headscarf use.[13] These reports indicate that the aggregate share of women wearing headscarves remained largely stable over the relevant period, remaining at approximately 61% in 2013 and 62% in 2018 (KONDA, 2019). As shown in Figure 1, aggregate share of veiled women exhibit only minor fluctuations across years, which supports the assumption that the relationship between observable characteristics and veiling behavior is unlikely to have changed meaningfully between the two DHS waves.

The prediction model includes a rich set of demographic and socioeconomic characteristics that are observed in both DHS waves. Importantly, we exclude measures of educational attainment from the set of predictors, as these variables are the main outcomes of interest in our analysis. We estimate the model using repeated cross-validation within the 2013 sample. In each iteration, we randomly reserve 30% of the observations for which veiling status is observed as a holdout set and train the model on the remaining 70%. Predictive performance is then evaluated on the holdout sample. Across 50 random splits, the random

[13] KONDA is a Turkish research company specializing in public opinion polling and social research. Its *KONDA Barometer* is a nationally representative monthly face-to-face surveys conducted since March 2010 that tracks political preferences, social attitudes, and economic conditions. To our knowledge, it is the only dataset that provides consistent time-series evidence on veiling rates in Turkey.

Figure 1: Coverage status among women

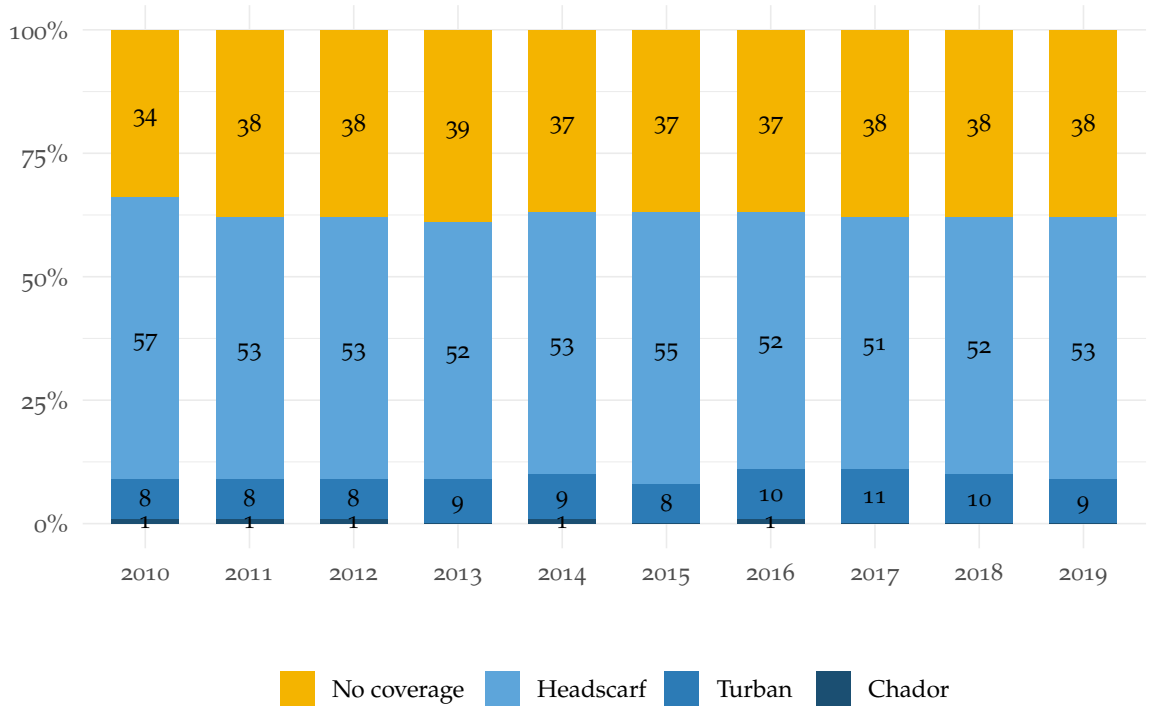

Notes: [Source: KONDA (2019)]

forest achieves an average out-of-sample AUC of 0.87 (sd = 0.007), implying that the model assigns a higher predicted probability of veiling to a randomly chosen veiled woman than to a randomly chosen unveiled woman in approximately 87% of pairwise comparisons. The corresponding classification accuracy is 0.81 (sd = 0.007), meaning that roughly 81% of women in the holdout sample are correctly classified using a 0.5 probability threshold. [14]

To account for some potential distributional shifts in the predictors between 2013 and 2018, we recalibrate the 2013-trained model's predicted probabilities so that their mean matches the observed aggregate veiling rate in the 2018 wave. Specifically, letting p_i denote the raw predicted probability, we apply an intercept shift on the log-odds scale:

$$p_i^{cal} = logit^{-1}\left(logit(p_i) + \delta\right),$$

where δ is chosen such that $\frac{1}{N}\sum_i p_i^{cal} = \bar{y}_{2018}$ (in our case $\delta = 0.58$). This calibration adjusts the level of predicted probabilities to be consistent with the 2018 prevalence while

[14] The AUC (Area Under the Receiver Operating Characteristic Curve) measures the model's ability to discriminate between veiled and unveiled women and ranges from 0.5 (no discriminatory power) to 1 (perfect discrimination). Classification accuracy refers to the share of correctly classified observations when predicted probabilities are converted into binary outcomes using a 0.5 threshold.

preserving the ranking of individuals.[15] Importantly, this quality leads to the metrics that depend only on the ranking (e.g., AUC) to remain unchanged.

To further illustrate the model's predictive performance, Figure 2 plots the distribution of predicted veiling probabilities in the 2013 sample separately by women's actual veiling status. Columns (4–6) of Table 1 present descriptive statistics for women classified as veiled or unveiled in 2018, using the observed 2018 veiling rate (0.62) as the cutoff for predicted probabilities. Additional descriptives on predictive accuracy is provided in Appendix A. Variable importance measures indicate that the strongest predictors of veiling behavior include indicators of media exposure such as internet use, region of residence, fertility and family formation patterns, and intergenerational education. Attitudinal measures related to gender norms also rank among the most important predictors.

4.2 Identification strategy

Our empirical strategy to identify effects of the removal of the veiling ban on women's schooling outcomes relies on a difference-in-differences (DiD) framework with two sources of variation. The first exploits cross-sectional differences in women's predicted likelihood of veiling, and the second exploits variation across birth cohorts, comparing women who had already passed university entrance age (18) at the time of the reform to those who were still below this age and therefore exposed to the policy change when making their university attendance decisions. As discussed in Section 2.1, the veiling ban was removed in 2010. This implies that cohorts born in 1992 or later were young enough for their university attendance decisions to be made under the post-reform regime. Therefore, the educational choices of women with a higher predicted likelihood of veiling who were born in or after 1992 can be interpreted as being exposed to the removal of the ban at the time relevant schooling decisions were taken. We estimate the following model to measure the effects on women's schooling outcomes:

$$Y_{it} = \beta_0 + \beta_1 \left(Veiled_i \times Post_t\right) + \beta_2 Veiled_i + FE_{\text{cohort}} + FE_{\text{province}} + \epsilon_{it}, \tag{1}$$

where Y_{it} denotes the educational outcome of individual i born in cohort t. $Post_t$ is an indicator equal to one for cohorts born in 1992 or later, and zero otherwise. $Veiled_i$ captures individual i's predicted likelihood of veiling. Cohort and province fixed-effects absorb common cohort trends and time-invariant regional heterogeneity. ϵ_{it} is the error

[15] This is because the mapping is strictly increasing in p_i.

term clustered at the province level unless specified otherwise.

The parameter of interest is β_1 which captures the differential change in educational outcomes for women with a higher likelihood of veiling in cohorts exposed to the removal of the ban, relative to older cohorts whose schooling decisions were completed prior to the reform. We restrict the sample to cohorts born between 1987-1998 in order to avoid conflating the effects of the veiling reforms with other schooling reforms taking place around the same time. [16]

It is important to note that since the veiling status is not directly observed but instead predicted, the classification is imperfect. Consequently, some women who actually veil may be assigned a low predicted probability (false-negative), while some who do not veil may be assigned a high predicted probability (false positive).[17] This imperfect treatment assignment which arises on both margins attenuates the estimated effects, implying that our coefficients should be interpreted as intention-to-treat (ITT) effects rather than treatment-on-the-treated effects.

Table 2 reports the results from estimating equation (1) on the years of schooling as the outcome variable. Column (1) uses the normalized predicted probabilities of veiling as a continuous measure of treatment intensity, so the higher values correspond to a greater likelihood of being veiled. [18] Column (2) instead uses a binary indicator that classifies women as veiled if their predicted probability exceeds the threshold of 0.62. [19] While using a probability cutoff provides a clear treatment definition, it imposes a single threshold and therefore fixes one realization of veiling status despite the underlying classification uncertainty. To account for this uncertainty, we implement a simulation-based multiple imputation (MI) procedure following Rubin, 1987. Specifically, in each iteration we draw binary veiling status for each individual from a Bernoulli distribution with success probability equal to the predicted probability of veiling, estimate equation (1) in the resulting imputed dataset, and combine the estimates using Rubin's rules. We repeat this procedure 100 times. The resulting pooled estimates are reported in Column (3) of Table 2. Figure B5 plots the distribution of the coefficients from this exercise.

We next estimate an event-study specification that allows the effects of the removal of

[16] Most notably, mandatory schooling was extended from 5 to 8 years in 1997, and further to 12 years in 2012.

[17] In the holdout sample in 2013, when using a threshold of 0.62, the model correctly classifies 75.1% of veiled women (1,439 out of 1,915) and 83.3% of unveiled women (839 out of 1,007). The corresponding false negative and false positive rates are 24.9% and 16.7%, respectively.

[18] Before normalization, the predicted probabilities of veiling in the 2018 sample have a mean of 0.68 and a standard deviation of 0.24.

[19] This threshold is chosen to match the average share of veiled women in 2018 reported by KONDA, 2019. The sensitivity of the estimates to this threshold choice is examined in Figure B6.

the veiling ban to vary across birth cohorts relative to the cohort that was 18 years old at the time of the reform. This exercise provides a test of the parallel trends assumption and allows us to examine how the reform's impact evolves across cohorts. Figure 3 presents the event-study results. The sample in the left panel (Panel a) includes the full sample of women, while the right panel (Panel b) restricts the analysis to non-migrants, i.e., those women who at the time of the survey lived in the same province in which they were born in order to avoid confounding effects of migration decisions taken later in life, which may themselves be shaped by schooling outcomes (Aydemir et al., 2021).

While the removal of the veiling ban may in principle affect schooling decisions more broadly, the policy is most likely to influence educational attainment at the university margin. Prior to the reform, women who chose to wear the veil faced restrictions in accessing higher education, while the veiling regulations in primary and secondary schooling remained unchanged and veiling remained prohibited throughout the study period. We therefore expect the reform's effects to be most pronounced for university attainment. To examine where in the schooling distribution these effects arise, we estimate equation (1) using a series of binary indicators for completing at least X years of schooling as outcome variables. Figure 4 plots the results from this exercise. Panel (a) uses the full sample, while Panel (b) restricts the analysis to non-migrants.

4.2.1 Alternative exposure measure

As an alternative way to capture exposure to the veiling ban, we exploit variation in pre-reform veiling prevalence across Turkey's 81 provinces. Provinces with higher veiling rates prior to the reform likely had a larger share of women whose access to higher education might have been constrained by the ban. The removal of the ban should therefore lead to larger educational responses in these provinces. To assess this, we estimate the following specification:

$$Y_{ipt} = \alpha_0 + \sum_{q=2}^{5} \alpha_q \left(Q_{pq} \times Post_t\right) + FE_{cohort} + FE_{birth\_province} + \varepsilon_{ipt}, \tag{2}$$

where Q_{pq} indicates that province p belongs to veiling quintile q, with the lowest quintile omitted, and Y_{ipt} is the outcome of interest for individual i born in province p and cohort t. We rely on veiling prevalence in individuals' birth provinces rather than their province of residence at the time of the survey to ensure that treatment reflects the social environment in which schooling decisions were likely formed, rather than the provinces to which they

migrated, which may be endogenous to prior schooling choices.

We construct province-level veiling rates using pre-reform DHS surveys that contain information on veiling behavior. Specifically, we pool the 2003 and 2008 DHS rounds and compute the share of veiled women in each province. Although the Turkish DHS is not designed to be representative at the province level, pooling the 2003 and 2008 surveys allows us to construct a useful proxy for cross-provincial differences in veiling prevalence prior to the reform. Figure 5 shows the resulting distribution of pre-reform veiling prevalence across provinces. Figure 6 presents the estimates from equation (2), and the corresponding regression results are reported in Table B1.

4.3 Robustness and validation

In what follows, we first evaluate the robustness of the veiling predictions that underpin our treatment measures. We then assess the stability of the difference-in-differences estimates across alternative treatment constructions, sample restrictions, and proxies for heterogeneity in the enforcement of the ban.

Alternative veil prediction model. We estimate an alternative, parametric measure of veiling propensity using a probit model. Specifically, in the 2013 DHS , where veiling status is directly observable, we model the probability that woman i is veiled as a function of her other observable characteristics $\mathbf{X}_i$: [20]

$$Pr(\text{veil}_i = 1 \mid X_i) = \Phi\big(\beta_0 + X_i^\top \beta\big) \,,$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. After estimating this model in the 2013 sample, we use the fitted coefficients to generate predicted veiling probabilities for the 2018 sample. Figure A2 presents the correlation between the two sets of probabilities estimated using this method with the ones estimated through a random forest. As shown in the figure, the predicted probabilities from the two approaches are highly correlated which reassures us that the predicted probabilities are not sensitive to the choice of prediction model. At the same time, the random forest achieves better separation between veiled and unveiled in the 2013 sample (Figure A3). We therefore retain the random-forest-based predictions in the main analysis while using the probit model as a

[20] As before, we exclude variables that directly capture one's educational attainment and retain only those covariates that are measured consistently across the two survey waves of 2013 and 2018.

validation exercise.

Robustness to treatment measurement. Throughout the analysis so far, the threshold-based veiling measure classifies women as veiled whenever their predicted probability of veiling exceeds 0.62, corresponding to the aggregate veiling rate observed in 2018. Although this choice provides a straightforward and empirically grounded threshold, the conversion of predicted probabilities into a binary treatment indicator is inevitably somewhat arbitrary. We therefore assess the robustness of the results to alternative cutoff values by re-estimating equation (1) across a grid of thresholds. Figure B6 reports the corresponding coefficient estimates and confidence intervals. Reassuringly, the estimates remain positive and broadly stable over a wide range of cutoffs which suggests that the main results do not hinge on the specific threshold used in the baseline analysis.

The university expansion program and access to universities. Among the other education reforms implemented in Turkey during this period, an important one was the expansion of higher education capacity through the establishment of 41 new universities between 2006 and 2008 (Caner et al., 2024). Figure B9 plots the number of universities over time and shows that this expansion took place shortly before the removal of the veiling ban. This expansion raises a concern for our interpretation. In provinces that did not have a university, women's educational attainment may have increased simply because local access to higher education improved, rather than the removal of the veiling ban per se. To address this concern, we re-estimate our main specifications after excluding provinces that did not have a university prior to 2006, i.e., before the expansion began. [21] This restriction allows us to examine whether our results are robust to removing the provinces most directly affected by the expansion in university supply. Table B2 presents the results for this exercise. The results remain similar, which helps rule out the concern that our main findings are driven by differential improvements in local access to higher education.

Pre-reform AKP support. As noted earlier, enforcement of the ban likely varied across universities, although we are unable to observe this variation directly due to data limitations. We use pre-reform provincial support for the AKP as a proxy for this heterogeneity, on the premise that enforcement may have been weaker in places where support for the party was stronger, given its longstanding opposition to the ban. We therefore divide provinces

[21] Before the expansion program began, 26 provinces did not have a university in 2006. By 2008, all but one province had at least one university. https://www.yok.gov.tr/

into quintiles based on AKP vote-shares in the general election in 2007 (Figure B8), [22] and re-estimate equation (1) separately within each group. Figure B7 presents the resulting estimates. The effects are broadly similar across the distribution. Given the similarity of the estimates across groups, it seems unlikely that cross-provincial variation in ban enforcement is driving the pattern we document.

Validation using observed veiling status. As a final robustness check, we re-estimate equation (1) using the 2013 DHS wave, in which veiling status is directly observed rather than predicted. Two caveats are worth noting. First, the younger cohorts in our sample, namely those born between 1992 and 1998, were between 15 and 21 years old at the time of the 2013 survey and may therefore not yet have completed their schooling trajectories. Second, the 2013 survey was conducted only three years after the removal of the veiling ban, which may be too soon for the full effects of the reform on completed educational attainment to be observed. With these caveats in mind, Table B3 reports the results. The estimated coefficient on the interaction term is positive, and consistent with the main results in Table 2, suggesting that the findings are not an artifact of the prediction procedure and that the underlying pattern is present in the data when veiling is directly observed.

5 Results

Table 2 reports the estimated effect of the removal of the ban on women's years of schooling. Columns (1)-(3) present the estimates for when we use the normalized predicted probability of veiling as a continuous measure for women's likelihood to be veiling, a binary indicator for whether one's predicted probability exceeds the threshold of 0.62, and a simulation-based multiple imputation procedure that assigns veiling status based on predicted probabilities, respectively. All columns include province and cohort fixed effects, and standard errors are clustered at the province level. Across all three columns, the interaction between being born post-1992 and the veiling measure is positive which suggests that the reform overall increased educational attainment for women more likely to be veiling.

In Column (1), with the normalized continuous veiling probability, the coefficient is 0.627 ($p < 0.05$), which indicates that for every one-standard-deviation increase in the predicted probability of veiling, women in the post-1992 cohorts gained roughly 0.63

[22] This is the last general election held before the 2010 reform.

additional years of schooling compared to their counterparts (who were too old at the time of the reform). The magnitude of the coefficient in Column (2) where a binary indicator for veiling is used is substantially larger, at 1.474 ($p < 0.01$). This estimate should be interpreted with caution, however, since the binary classification treats veiling status as known with certainty and does not account for uncertainty in the predicted probabilities. Column (3) addresses this concern by incorporating uncertainty in veiling assignment through simulation-based multiple imputation. Reassuringly, the resulting coefficient remains positive and statistically significant at the 90% level, 0.877 ($p < 0.1$).

Figure B6 further shows that the positive effect in the threshold-based veiling measure is not driven by the particular cutoff choice. Across a wide range of thresholds, the estimated coefficient remains positive and generally statistically significant.

The same pattern emerges when exposure to the reform is measured using an alternative proxy based on pre-reform veiling prevalence at the province level, rather than individual predicted veiling probabilities. Figure 6 shows that the gains in schooling are concentrated among women born in provinces in top quintile of pre-reform veiling rates. Across specifications, the estimated effect for quintile five is large and statistically significant, while coefficients for lower quintiles are modest and generally not distinguishable from zero.

Next, we examine the cohort-specific effects of the reform in an event-study framework that again uses predicted veiling probabilities to measure exposure. Figure 3 plots the coefficients from regressions that interact cohort indicators with veiling measures, omitting the 1992 cohort as the reference group. The figure confirms that for the cohorts who were not exposed to the reform, the coefficients are generally close to zero and do not show differential trends with veiling likelihood. For cohorts born after 1992, the estimates are positive and increase steadily. This pattern is visible under both the continuous probability measure (the blue series) and the threshold-based binary one (the orange series). Here, again the binary measure displays somewhat larger estimates, consistent with the results in Table 2. The similarity of the estimates in Panel b, where the sample is restricted to women who reside in the same province in which they were born, suggests that the main results are not driven by selective migration across provinces. [23]

Having established that years of schooling increased following the reform for women who are more likely to be veiling, we next turn to examine which margins of educational attainment are most affected by the reform. To do so, we re-estimate equation (1) using

[23] In our sample, 27.6% of women reside in a province different from their province of birth. Pre-trends tests pass for both veiling measures in the non-migrant sample, in the baseline sample, the continuous veiling measure shows some evidence of pre-existing trends, while the binary threshold-based measure does not.

indicators for having completed at least X years of schooling, where $X \in \{8, \ldots, 14\}$ as outcomes. Figure 4 shows that the positive effects are somewhat concentrated around the transition into and within secondary schooling, i.e., 8 to 10 years of schooling, while they become smaller at higher schooling thresholds. The same broad pattern is visible across the three veiling measures although the magnitudes differ. The threshold-based measure tends to yield the largest point estimates, the continuous probability measure produces more modest coefficients, and the multiple-imputation estimates usually lie in between and have wider confidence intervals due to the additional uncertainty they incorporate. Once again, the non-migrant sample (Panel b) yields a very similar pattern.

6 Discussion

Our findings indicate that the removal of the veiling ban improved veiled women's educational attainment in Turkey, with the gains being concentrated around secondary schooling. This suggests that a backlash effect, as predicted by Carvalho et al., 2024, was at work, leading girls to drop out of school in response to the ban. Importantly, these gains accrued in levels of education where veiling was banned throughout the study period for the cohorts in our sample, indicating that even symbolic policy changes that increase acceptance of religious expression in the schooling system may attenuate the backlash effect.

The impacts of the removal of the ban were concentrated among the younger cohorts in our sample, in the most conservative provinces. Previous literature has shown that such anticipatory human capital investments can arise from expanded future opportunities (see Sandholtz, 2024 for an overview), as well as from increases in educational aspirations (in our case - higher educational aspirations for veiled women). A natural question that follows is why these gains in schooling do not translate into gains in tertiary education. One possible explanation is that uneven enforcement of the veiling ban attenuates the impact of its removal, similar to a key mechanism in Meyersson (2014). Although we cannot directly observe enforcement, we can tentatively rule out substantial heterogeneity by enforcement intensity, as treatment effects do not differ between provinces where enforcement would be expected to be more or less strict.

An alternative interpretation consistent with these results is that the effects are concentrated among the most constrained women in our sample, i.e., those who would otherwise have dropped out of the education system after completing the mandatory eight years of

schooling, well before tertiary education. This may reflect a bundle of constraints discussed in Section 2. While the removal of the ban encourages them to remain in the schooling system longer, it does not induce progression to higher levels.

While the limited sample sizes in our data do not permit a more detailed investigation of additional mechanisms, future research can examine the constraints faced by veiled women in accessing higher education in greater detail.

References

Abdelgadir, A. and V. Fouka (2020). "Political secularism and Muslim integration in the West: Assessing the effects of the French headscarf ban". In: *American Political Science Review* 114.3, pp. 707–723.

Akerlof, G. A. and R. E. Kranton (2000). "Economics and identity". In: *The Quarterly Journal of Economics* 115.3, pp. 715–753.

Akerlof, G. A. and R. E. Kranton (2002). "Identity and schooling: Some lessons for the economics of education". In: *Journal of Economic Literature* 40.4, pp. 1167–1201.

Albayrak, H. Ş. (2023). "Justification of headscarf bans in Europe and Turkey: A comparative analysis". In: *Asian Journal of Women's Studies* 29.4, pp. 484–505.

Appadurai, A. (2004). "The Capacity to Aspire: Culture and the Terms of Recognition". In: *Culture and Public Action*. Ed. by V. Rao and M. Walton. Stanford, CA: Stanford University Press, pp. 59–84.

Aydemir, A., M. G. Kirdar, and H. Torun (2021). *The effect of education on geographic mobility: Incidence, timing, and type of migration*. Tech. rep. IZA Discussion Papers.

Balci, B. and N. Monceau (2023). *Turkey, a Century of Change in State and Society*. Springer Nature.

Bisin, A. and T. Verdier (2000). "'"Beyond the melting pot": cultural transmission, marriage, and the evolution of ethnic and religious traits". In: *The Quarterly Journal of Economics* 115.3, pp. 955–988.

Bisin, A. and T. Verdier (2001). "The economics of cultural transmission and the dynamics of preferences". In: *Journal of Economic Theory* 97.2, pp. 298–319.

Bursztyn, L., G. Egorov, and R. Jensen (2019). "Cool to be smart or smart to be cool? Understanding peer pressure in education". In: *The Review of Economic Studies* 86.4, pp. 1487–1526.

Bursztyn, L., A. L. González, and D. Yanagizawa-Drott (2020). "Misperceived Social Norms: Women Working Outside The Home in Saudi Arabia". In: *American Economic Review* 110.10, pp. 2997–3029.

Caner, A., M. Derebasoglu, and C. Okten (2024). "Attainment and gender equality in higher education: evidence from a large-scale expansion". In: *Journal of Human Capital* 18.3, pp. 469–530.

Çarkoğlu, A. and B. Toprak (2007). *Religion, society and politics in a changing Turkey*. Tesev.

Carvalho, J.-P. (2013). "Veiling". In: *The Quarterly Journal of Economics* 128.1, pp. 337–370.

Carvalho, J.-P., M. Koyama, and C. Williams (2024). "Resisting Education". In: *Journal of the European Economic Association* 22.6, pp. 2549–2597.

Cindoğlu, D. (2011). *Headscarf ban and discrimination: Professional headscarved women in the labor market*. TESEV publications İstanbul, Turkey.

Corekcioglu, G. (2021). "Unveiling the effects of a headscarf ban: Evidence from municipal jobs in Turkey". In: *Journal of Comparative Economics* 49.2, pp. 382–404.

Dalton, P. S., S. Ghosal, and A. Mani (2016). "Poverty and aspirations failure". In: *The Economic Journal* 126.590, pp. 165–188.

Delavande, A. and B. Zafar (2019). "University choice: The role of expected earnings, nonpecuniary outcomes, and financial constraints". In: *Journal of Political Economy* 127.5, pp. 2343–2393.

Dhar, D., T. Jain, and S. Jayachandran (2022). "Reshaping Adolescents' Gender Attitudes: Evidence from a School-based Experiment in India". In: *American Economic Review* 112.3, pp. 899–927.

Genicot, G. and D. Ray (2017). "Aspirations and inequality". In: *Econometrica* 85.2, pp. 489–519.

Gulesci, S. and E. Meyersson (2015). *For the Love of the Republic: Education, Secularism, and Empowerment*. Working Paper 490. Bocconi University, Innocenzo Gasparini Institute for Economic Research.

Kökkizil, M. (2022). *Parental Religiosity and Missing School-Girls in Turkey*. Tech. rep. Faculty of Economics and Management at the Free University of Bozen.

KONDA (2019). *Barometer: Political and Social Survey Series, 100th Issue*. Survey Report. KONDA Research and Consultancy.

Maurin, É. and N. Navarrete H (2023). "Behind the veil: the effect of banning the Islamic veil in schools". In: *Economic Policy* 38.113, pp. 63–98.

Meyersson, E. (2014). "Islamic Rule and the Empowerment of the Poor and Pious". In: *Econometrica* 82.1, pp. 229–269.

Oh, S. (2023). "Does identity affect labor supply?". In: *American Economic Review* 113.8, pp. 2055–2083.

Rubin, D. B. (1987). *Multiple Imputation for Nonresponse in Surveys*. New York: John Wiley Sons.

Sandholtz, W. A. (2024). *Secondary School Access Raises Primary School Achievement*. CESifo Working Paper. CESifo.

Shofia, N. (2022). "Why Veil? Religious Headscarves and the Public Role of Women". Unpublished manuscript, July 21, 2022.

Uğur, Z. B. (2020). "Unveiled: The effect of the headscarf ban on women's tertiary education in Turkey". In: *Feminist Economics* 26.2, pp. 187–217.

Yurdakul, E. and W. Lu (2025). "To Veil or Not to Veil? Assessing the Removal of Headscarf Ban in a Muslim Country". Unpublished manuscript.

Zurcher, E. J. (2004). *Turkey: A Modern History*. London: I.B. Tauris, p. 424.

Figures and Tables

Table 1: Summary statistics by veiling status

	DHS 2013 (observed)			DHS 2018 (predicted)		
	Unveiled mean (SD) (1)	*Veiled* mean (SD) (2)	*Diff* (1)-(2) (3)	*Unveiled* mean (SD) (4)	*Veiled* mean (SD) (5)	*Diff* (4)-(5) (6)
Demographics						
Age	28.54 (9.74)	32.57 (9.52)	−4.04***	29.61 (9.69)	34.17 (9.47)	−4.56***
Married	0.52 (0.50)	0.80 (0.40)	−0.28***	0.55 (0.50)	0.83 (0.37)	−0.28***
Household size	4.33 (2.01)	5.46 (2.60)	−1.13***	3.93 (1.39)	5.52 (2.16)	−1.59***
Rural	0.14 (0.35)	0.32 (0.46)	−0.17***	0.14 (0.34)	0.39 (0.39)	−0.26***
Mother literate	0.75 (0.43)	0.41 (0.49)	0.34***	0.84 (0.36)	0.33 (0.47)	0.51***
Educational outcomes						
Years of schooling	10.13 (3.63)	5.88 (3.94)	4.26***	11.20 (3.55)	5.73 (3.61)	5.46***
Primary or below	0.18 (0.38)	0.63 (0.48)	−0.45***	0.12 (0.33)	0.63 (0.48)	−0.51***
Secondary or above	0.52 (0.50)	0.15 (0.36)	0.36***	0.74 (0.44)	0.28 (0.45)	0.45***
Attitudes toward wife beating is justified if she						
Argues with the husband	0.03 (0.18)	0.10 (0.30)	−0.06***	0.01 (0.11)	0.07 (0.26)	−0.06***
Burns the food	0.01 (0.07)	0.02 (0.14)	−0.02***	0.00 (0.05)	0.02 (0.13)	−0.02***
Goes out without telling the husband	0.02 (0.13)	0.08 (0.27)	−0.06***	0.01 (0.09)	0.06 (0.24)	−0.05***
Neglects the children	0.05 (0.22)	0.13 (0.34)	−0.08***	0.02 (0.15)	0.11 (0.32)	−0.09***
Refuses sex with the husband	0.01 (0.09)	0.05 (0.21)	−0.04***	0.00 (0.06)	0.03 (0.18)	−0.03***
Observations	3,354	6,382		3,363	3,983	

Note: This table reports summary statistics by veiling status for women aged 15–49 in the DHS 2013 and DHS 2018 samples. In the 2013 sample, veiling status is directly observed in the survey. In the 2018 sample, veiling status is constructed using the predicted classification from the main prediction model using a threshold of 0.62. Columns (1), (2), (4), and (5) report group means, with standard deviations in parentheses. Columns (3) and (6) report the difference in means between unveiled and veiled women. The variables on attitudes capture the share of respondents who report that wife beating is justified under each listed circumstance.

Table 2: Effects of the removal of the veiling ban on women's years of schooling

	Dependent variable:		
	Years of schooling		
	(1) Probability	(2) Threshold	(3) Multiple Imputation
Post_92 × Veiling Measure	0.627**	1.474***	0.877*
	(0.254)	(0.405)	(0.526)
Birthyear FE	✓	✓	✓
Province FE	✓	✓	✓
Observations	2,452	2,452	2,452
MI iterations	—	—	100
R^2	0.560	0.468	—

Note: This table reports estimates of equation (1). The sample includes women born between 1987 and 1998 who were surveyed in the 2018 DHS round. Column (1) uses the normalized continuous predicted probability of veiling, column (2) uses a threshold-based binary veiling measure with the threshold set at 0.62, and column (3) uses multiple imputation to account for uncertainty in veiling status. The estimates in column (3) are pooled across 100 imputations. Post_92 is an indicator equal to one for cohorts born in 1992 or later. All specifications include birthyear and province fixed effects. Sampling weights are applied, and standard errors clustered at the province level are reported in parentheses. The mean years of schooling among pre–1992 cohorts is 8.96. *p<0.1; **p<0.05; ***p<0.01.

Figure 2: Distribution of predicted veiling probabilities by observed status

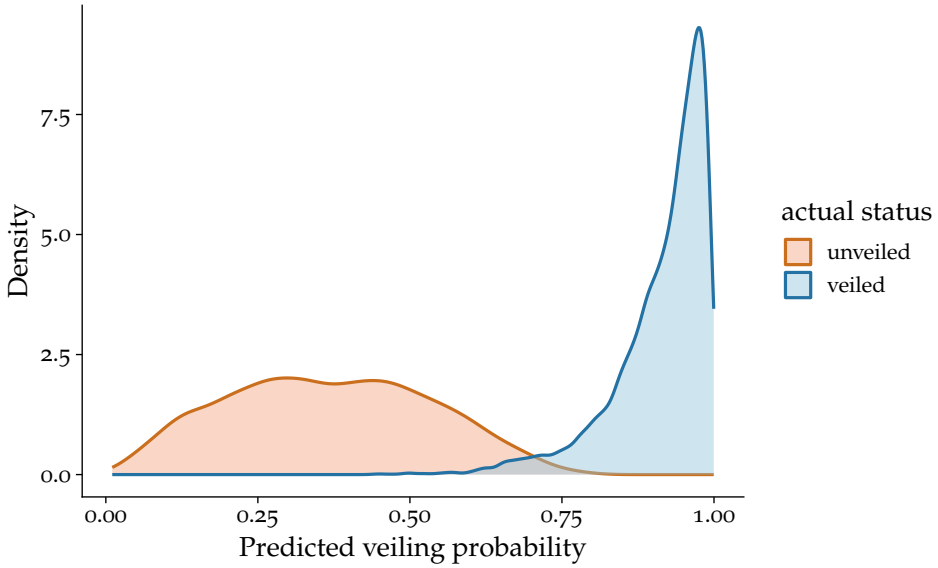

Notes: This figure plots the distribution of predicted probabilities from the random forest model separately for women who report being veiled and unveiled in the 2013 DHS.

Figure 3: Event-study estimates of the removal of the veiling ban

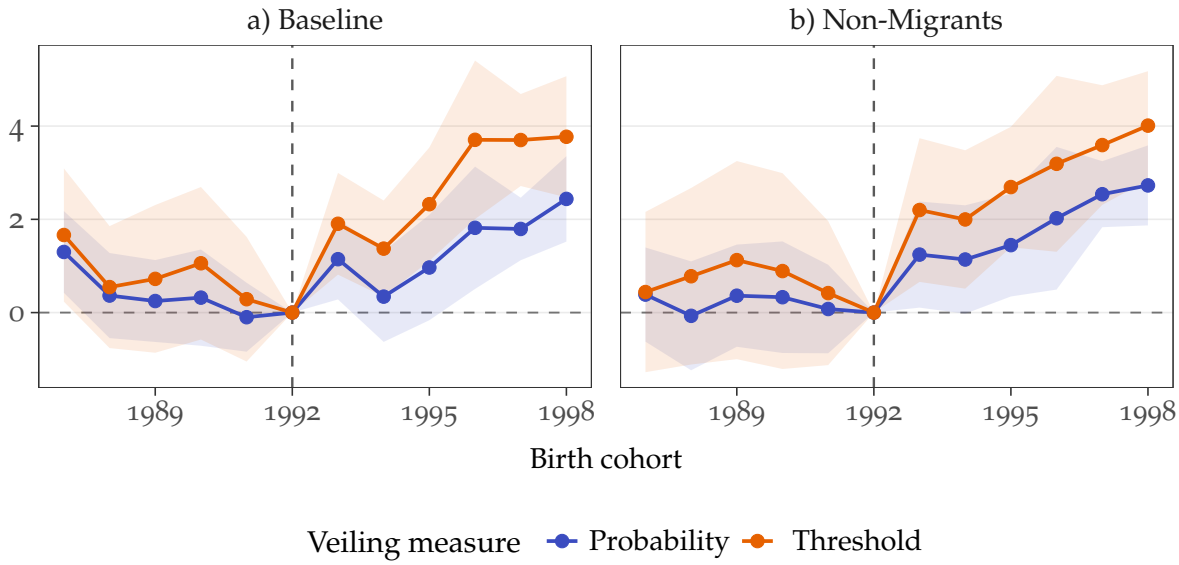

Notes: This figure plots the event-study estimates of the effect of the removal of the veiling ban on years of schooling across birth cohorts. Coefficients are obtained from regressions that interact cohort indicators with women's likelihood of veiling, omitting the 1992 birth cohort, which is the cohort aged 18 at the time of the reform. The sample includes women born in 1987-1998, and surveyed in the 2018 Turkish DHS wave. Panel (a) uses the full sample, while Panel (b) restricts the sample to non-migrants, defined as women residing in their childhood province at the time of the survey. The blue series uses the normalized predicted probability of veiling as a continuous treatment measure, while the orange series uses a binary indicator equal to one if the predicted probability exceeds the threshold of 0.62. Shaded areas denote 95% confidence intervals. All specifications include cohort and province fixed effects. Sampling weights are applied and standard errors are clustered at the province level.

Figure 4: Effects of the removal of the veiling ban across schooling thresholds

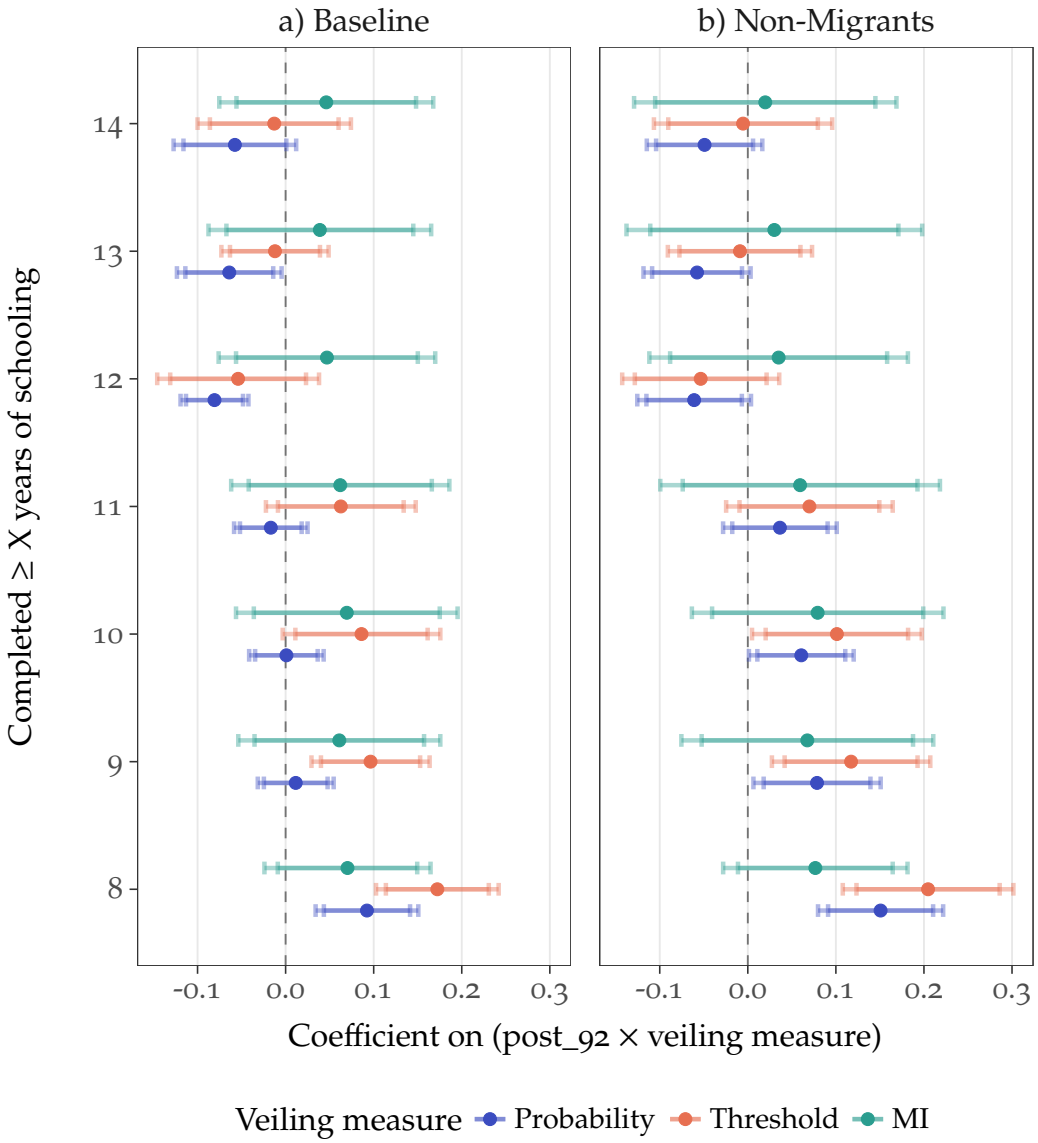

Notes: This figure reports estimates of equation (1) using binary indicators for completing at least X years of schooling as outcome variables, where $X \in \{8, \ldots, 14\}$. The sample includes women born in 1987-1998 in the 2018 Turkish DHS wave. Each point corresponds to the coefficient on the interaction between the post-1992 indicator and the veiling measure. Panel (a) reports estimates for the full sample, while Panel (b) restricts the sample to non-migrants, defined as women residing in their childhood province at the time of the survey. The blue series uses the normalized predicted probability of veiling as a continuous treatment measure, the orange series uses a binary indicator equal to one if the predicted probability exceeds 0.62, and the green series reports multiple-imputation estimates from 100 iterations. Horizontal bars denote 90% (inner) and 95% (outer) confidence intervals. All regressions include birth-year and childhood-province fixed effects. Sampling weights are applied and standard errors are clustered at the province level.

Figure 5: Pre-reform veiling prevalence across provinces

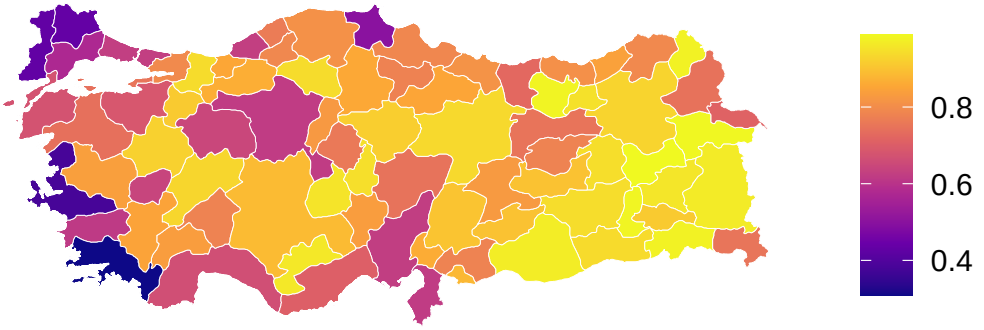

Notes: This map plots province-level veiling prevalence prior to the reform, measured as the share of women who report wearing a headscarf in the pooled 2003 and 2008 Turkish DHS rounds.

Figure 6: Effects of the removal of the veiling ban by pre-reform veiling quintile

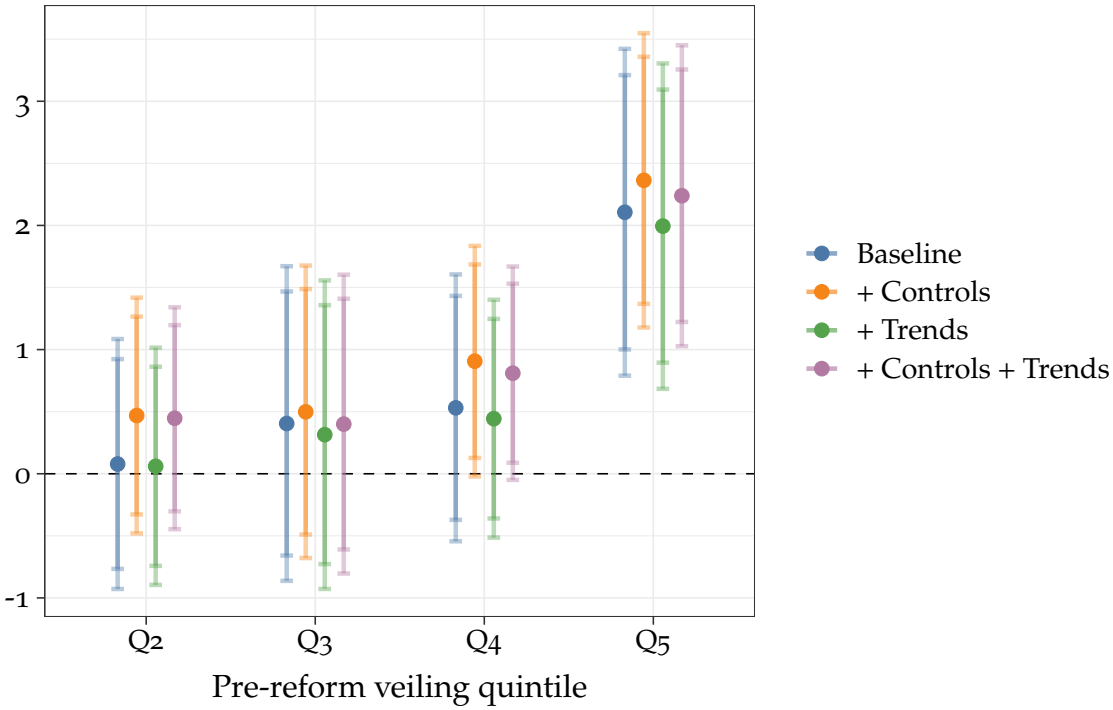

Notes: This figure reports estimates from equation (2) that interact the post-1992 indicator with quintiles of province-level pre-reform veiling prevalence. Quintiles are constructed using the share of veiled women in the pooled 2003 and 2008 Turkish DHS rounds. The omitted category is the lowest-veiling quintile (Q1), so each coefficient is interpreted relative to that group. Points correspond to coefficient estimates and bars denote 90% and 95% confidence intervals. The four series report the baseline specification, the specification with controls, the specification with province-specific linear trends, and the specification including both controls and province-specific trends. The corresponding regression results are reported in Table B1.

Appendix A Classification Performance and Validation

Figure A1: Predicted and observed veiling shares by probability decile

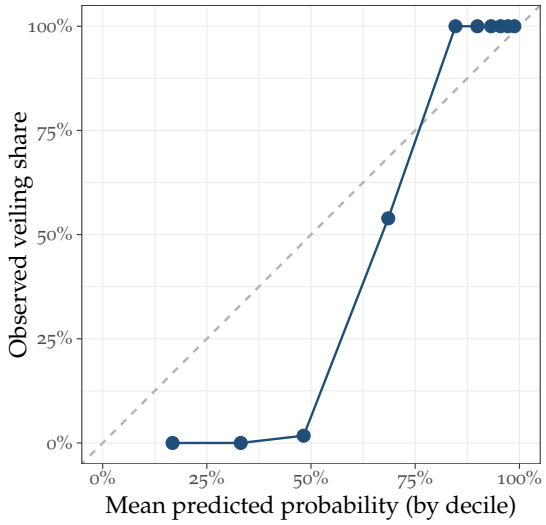

Notes: This figure plots observed veiling shares against mean predicted probabilities by decile of the predicted distribution in the 2013 DHS.

Figure A2: Correlation between predicted veiling probabilities

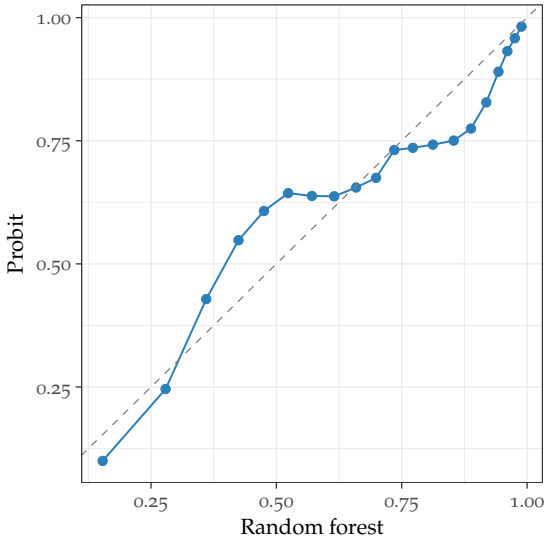

Notes: This figure compares individual-level predicted veiling probabilities from the random forest (x-axis) and the probit (y-axis) model. To reduce noise, observations are sorted by random-forest predictions and grouped into ventiles; each marker is the mean prediction within a bin. The dashed 45-degree line denotes equality across models. The close alignment indicates agreement between the two prediction methods, while deviations from the line show where one model is more or less extreme.

Figure A3: Distribution of predicted veiling probabilities using a probit model by observed status

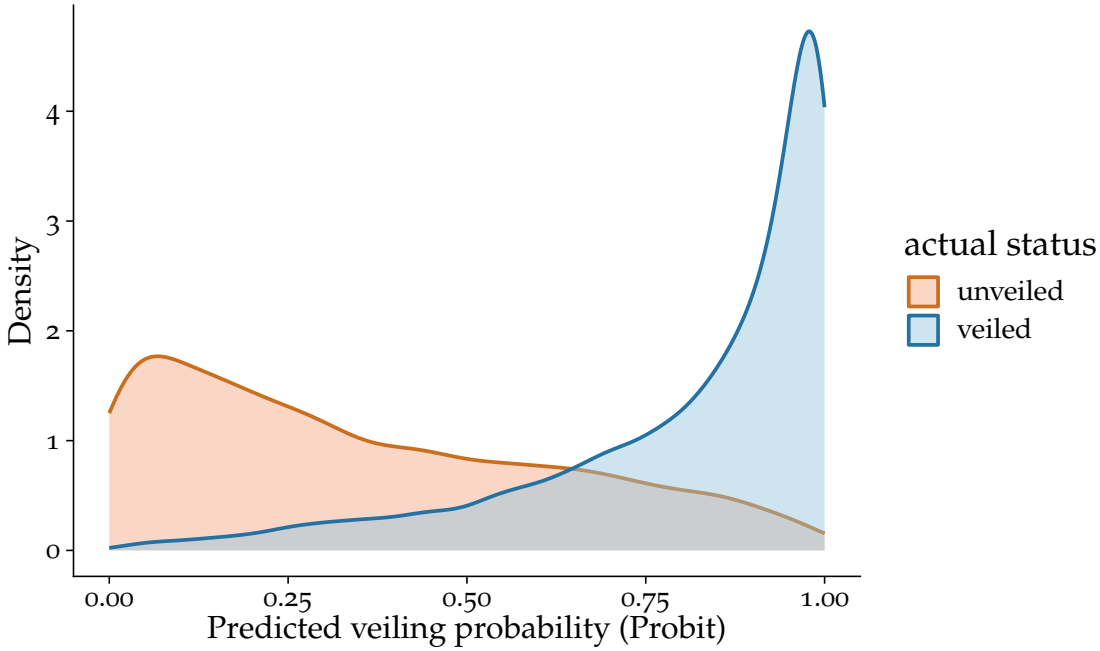

Notes: This figure plots the distribution of predicted probabilities using a probit model separately for women who report being veiled and unveiled in the 2013 DHS.

Appendix B Additional stuff

Figure B1: Years of schooling by age and veiling status

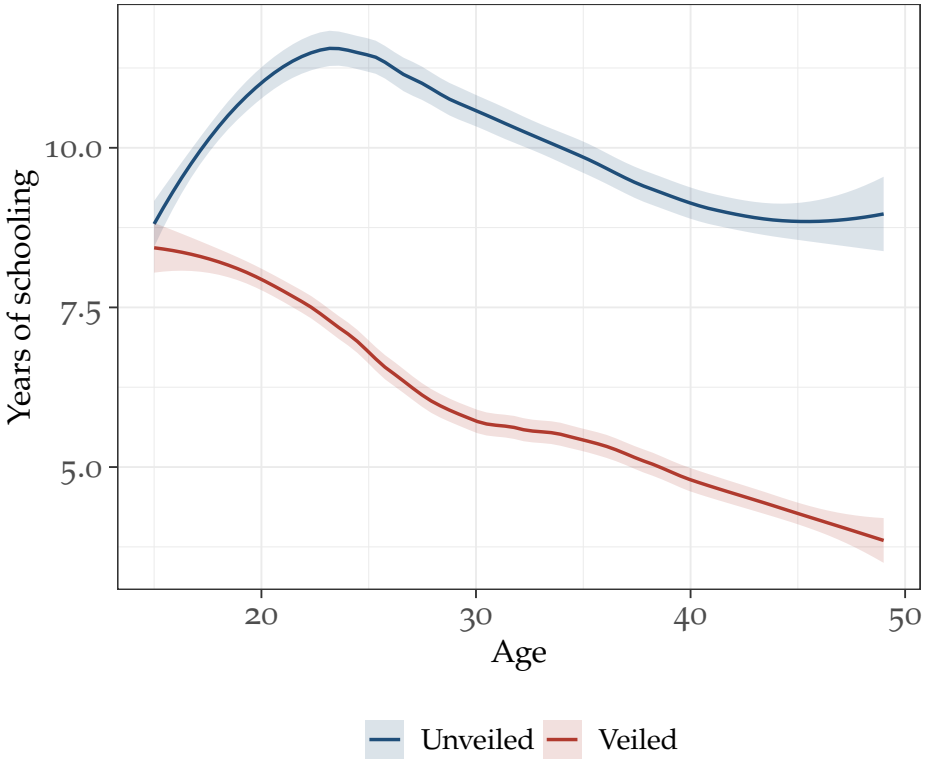

Notes: This figure plots mean years of schooling by age separately for veiled (red) and unveiled (blue) women. The sample includes women aged 15 to 49 surveyed in the 2013 round of the Turkish DHS. Shaded areas represent 95% confidence intervals.

Figure B2: Veiling rates by years of schooling

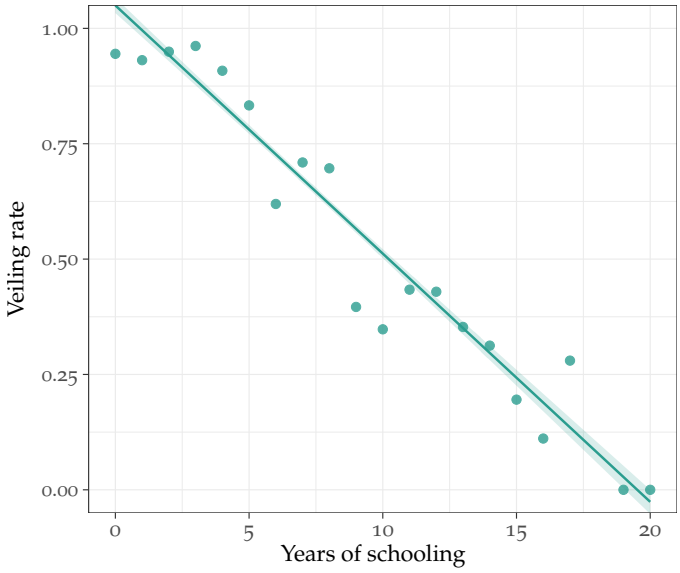

Notes: This figure plots veiling rates against years of schooling for women in the 2013 round of the Turkish DHS. Each point represents the mean veiling rate for a given level of schooling. The solid line shows the fitted linear relationship, and the shaded area represents the 95% confidence interval.

Figure B3: Veiling rates by age group and marital status

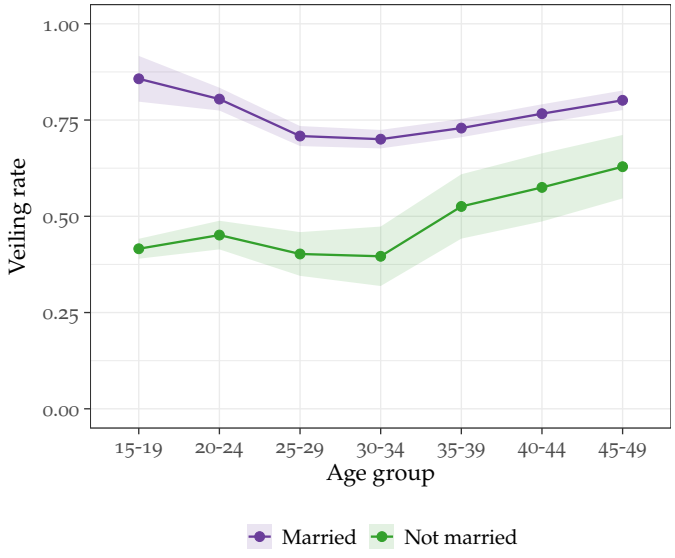

Notes: This figure plots veiling rates by five-year age group separately for married (purple) and unmarried (green) women. The sample includes women aged 15 to 49 surveyed in the 2013 round of the Turkish DHS. Shaded areas represent 95% confidence intervals.

Figure B4: Composition of veiled and unveiled women by attitude-index quintile

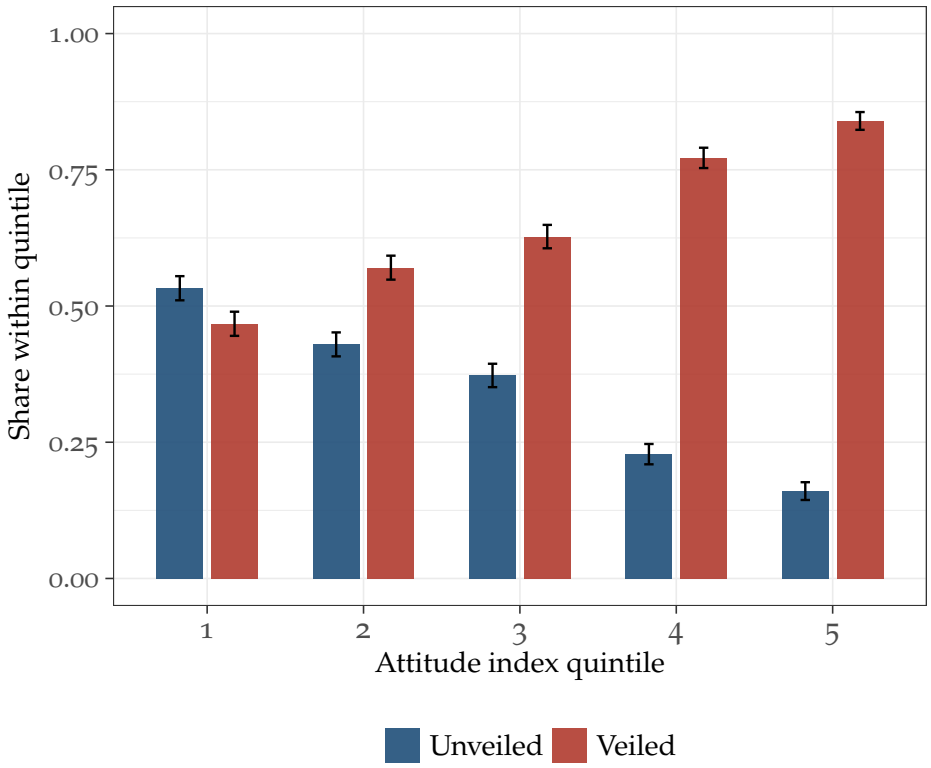

Notes: This figure groups women into five quintiles based on an attitude index, ordered from lower to higher acceptance of wife-beating in the following circumstances: if the wife refuses to have sex with her husband, neglects the children, goes out without telling her husband, argues with her husband, or burns the food. For each quintile, the figure shows the share of women who are veiled and unveiled in the 2013 round of the Turkish DHS. Error bars represent 95% confidence intervals.

Figure B5: Distribution of multiple-imputation estimates

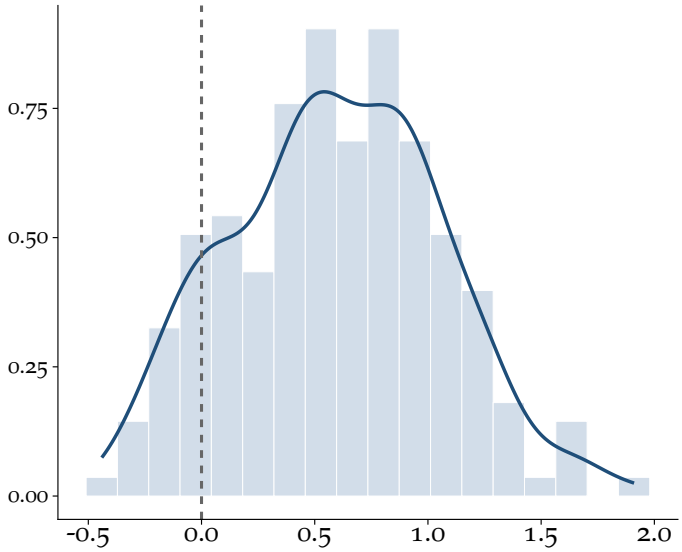

Notes: This figure plots the distribution of coefficient estimates obtained from 100 imputed datasets. In each iteration, binary veiling status is drawn for each individual from a Bernoulli distribution with success probability equal to the predicted probability of veiling, and equation (1) is then re-estimated on the resulting dataset. The histogram and kernel density summarize the distribution of the coefficient on the interaction between the post-1992 indicator and imputed veiling status. The dashed vertical line at zero is shown for reference.

Figure B6: Sensitivity of threshold-based estimates to the classification cutoff

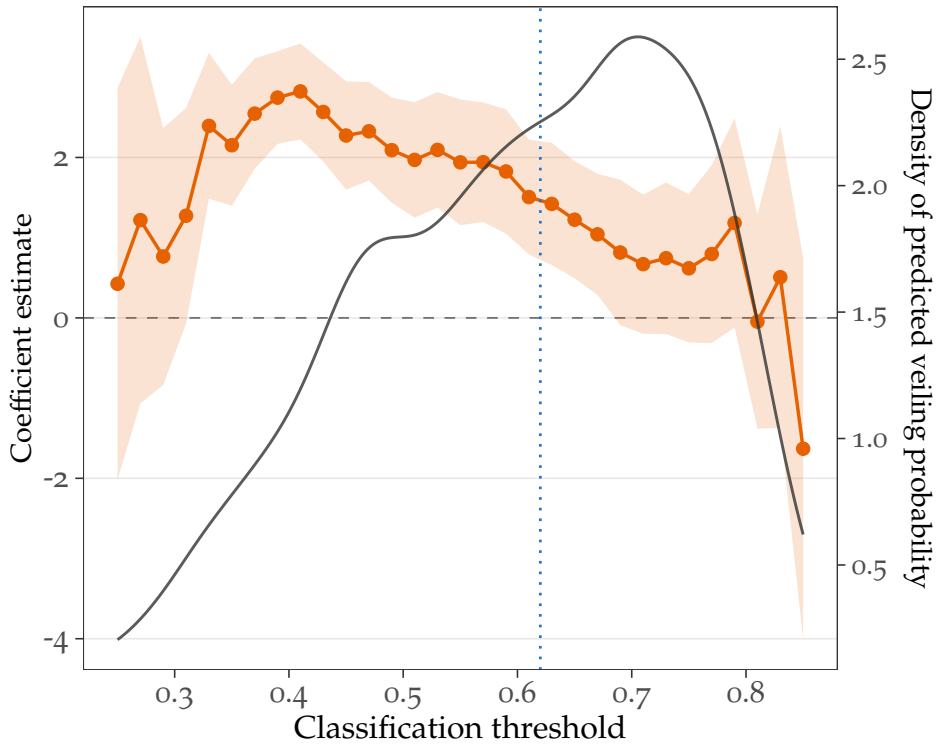

Notes: This figure plots how the estimated coefficient on the interaction between the post-1992 indicator and binary veiling status varies with the classification threshold used to define predicted veiling. For each threshold on the horizontal axis, women are classified as veiled if their predicted probability exceeds that cutoff, and equation (1) is re-estimated. The orange line reports the coefficient estimates and the shaded band denotes 95% confidence intervals. The gray line shows the distribution of predicted veiling probabilities in the sample. The vertical dotted line marks the observed aggregate veiling rate in 2018, equal to 0.62, which is used as the baseline threshold in the main analysis.

Figure B7: Heterogeneity by pre-reform AKP vote share

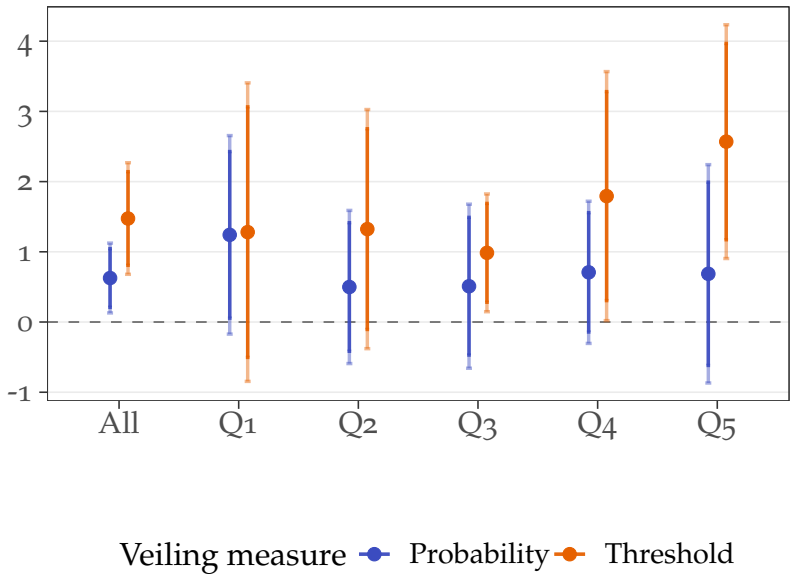

Notes: This figure reports estimates of equation (1) estimated separately for the full sample ("All") and within each AKP vote-share in 2007 quintile (Q1–Q5). The sample includes women born in 1987-1998 in the 2018 Turkish DHS wave. Province and birth-year fixed effects are included in all specifications. The blue series uses the normalized predicted probability of veiling as a continuous treatment measure; the orange series uses a binary indicator equal to one if the predicted probability exceeds 0.62. Whiskers show 95% (outer) and 90% (inner) confidence intervals. Sampling weights are applied and standard errors are clustered at the province level.

Figure B8: Provincial AKP vote-share quintiles (2007)

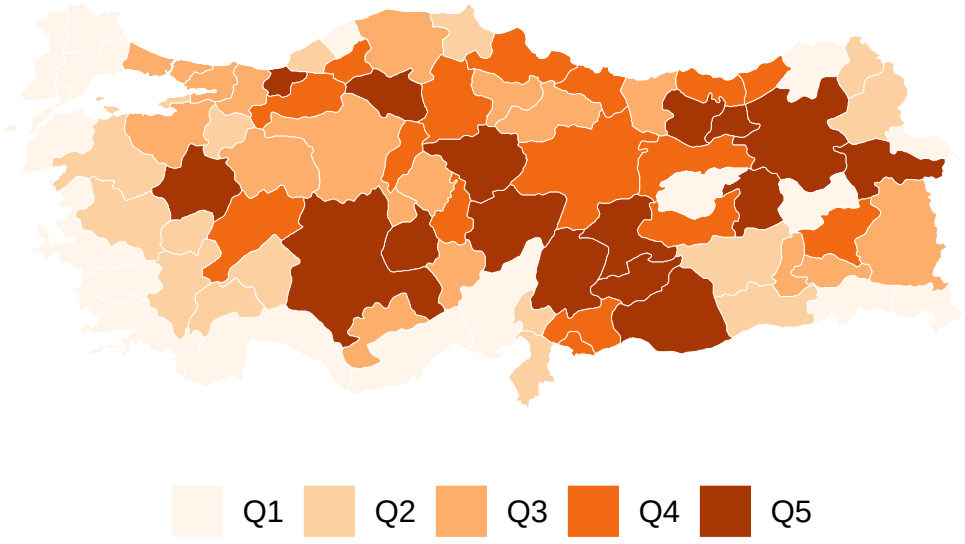

Notes: This map groups provinces into quintiles based on AKP vote share in 2007 general election (Q1 lowest to Q5 highest). Darker shading indicates higher AKP support.

Figure B9: Number of universities in Turkey over time

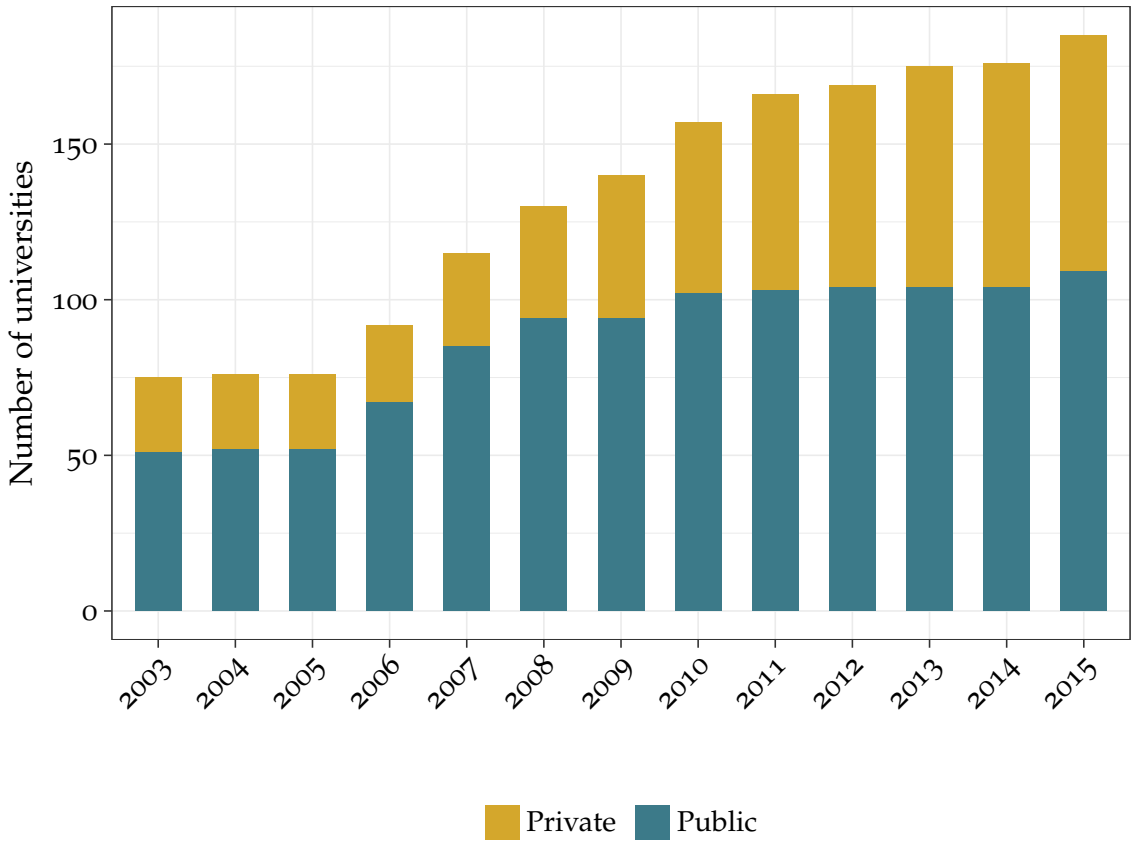

Notes: This figure plots the annual number of public and private universities in Turkey.

Table B1: Pre-reform veiling share: quintiles

	Dependent variable:			
	Years of schooling			
	(1)	(2)	(3)	(4)
Post_92 × Veiling Quintile 2	0.079	0.469	0.060	0.447
	(0.513)	(0.484)	(0.487)	(0.456)
Post_92 × Veiling Quintile 3	0.405	0.499	0.315	0.400
	(0.646)	(0.601)	(0.634)	(0.614)
Post_92 × Veiling Quintile 4	0.531	0.906*	0.444	0.810*
	(0.548)	(0.474)	(0.489)	(0.438)
Post_92 × Veiling Quintile 5	2.106***	2.363***	1.994***	2.239***
	(0.671)	(0.605)	(0.669)	(0.618)
Birthyear FE	✓	✓	✓	✓
Childhood province FE	✓	✓	✓	✓
Controls		✓		✓
Province-specific trends			✓	✓
Observations	2,388	2,388	2,388	2,388
R^2	0.272	0.412	0.272	0.413

Note: This table reports estimates of equation (2). The sample includes women born between 1987 and 1998 who were surveyed in the 2018 DHS round. Column (1) uses the normalized continuous predicted probability of veiling, column (2) uses a threshold-based binary veiling measure with the threshold set at 0.62, and column (3) uses multiple imputation to account for uncertainty in veiling status. The estimates in column (3) are pooled across 100 imputations. 'Post_92' is an indicator equal to one for cohorts born in 1992 or later. All specifications include birthyear and province fixed effects. Sampling weights are applied, and standard errors clustered at the province level are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table B2: Effects of the removal of the veiling ban on women's years of schooling, excluding provinces without a university in 2006

	Dependent variable:		
	Years of Schooling		
	(1) Probability	(2) Threshold	(3) Multiple Imputation
Post_92 × Veiling Measure	0.525*	1.184***	0.738*
	(0.295)	(0.432)	(0.591)
Birthyear FE	✓	✓	✓
Province FE	✓	✓	✓
Observations	2,033	2,033	2,033
MI iterations	—	—	100
R^2	0.550	0.456	—

Note: This table reports estimates of equation (1). The sample includes women born between 1987 and 1998 who were surveyed in the 2018 DHS round. 26 provinces that did not have university in 2006 i.e., before the university expansion program began are excluded from the sample. Column (1) uses the normalized continuous predicted probability of veiling, column (2) uses a threshold-based binary veiling measure with the threshold set at 0.62, and column (3) uses multiple imputation to account for uncertainty in veiling status. The estimates in column (3) are pooled across 100 imputations. Post_92 is an indicator equal to one for cohorts born in 1992 or later. All specifications include birthyear and province fixed effects. Sampling weights are applied, and standard errors clustered at the province level are reported in parentheses. The mean years of schooling among pre–1992 cohorts is 9.34. *p<0.1; **p<0.05; ***p<0.01.

Table B3: Validation using observed veiling status in the 2013 DHS

	Dependent variable:	
	Years of schooling	
	(1)	(2)
Post_92 × Actual Veiling Status	2.534***	2.101***
	(0.320)	(0.289)
Birthyear FE	✓	✓
Province FE	✓	✓
Controls		✓
Observations	3,476	3,476
R^2	0.332	0.420

Note: This table reports estimates of equation (2). The sample includes women born between 1987 and 1998 who were surveyed in the 2018 DHS round. Quintiles are constructed using the share of veiled women in the pooled 2003 and 2008 Turkish DHS rounds. The omitted category is the lowest-veiling quintile (Q1), so each coefficient is interpreted relative to that group. Column (1) reports the baseline specification. Column (2) adds controls for age, household size, urban status, and marital status. Column (3) includes province-specific linear trends. Column (4) includes both controls and province-specific trends. All specifications include birthyear and childhood province fixed effects. Sampling weights are applied, and standard errors clustered at the province level are reported in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.