Université Fédérale
Toulouse Midi-Pyrénées

THÈSE

En vue de l'obtention du
DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par l'Université Toulouse 1 Capitole

Présentée et soutenue par

Antoine JACQUET

Le 26 septembre 2023

Essays in the economics of culture

Ecole doctorale : **TSE - Toulouse Sciences Economiques**

Spécialité : **Sciences Economiques - Toulouse**

Unité de recherche :
TSE-R - Toulouse School of Economics - Recherche

Thèse dirigée par
Paul SEABRIGHT

Jury

Mme Sriya IYER, Rapporteure
M. Thierry VERDIER, Rapporteur
M. Augustin TAPSOBA, Examinateur
M. Alfred GALICHON, Examinateur
M. Paul SEABRIGHT, Directeur de thèse
Mme Emmanuelle AURIOL, Présidente

Essays in the economics of culture

Antoine Jacquet

Thesis submitted to

Toulouse School of Economics

in partial fulfilment for the award of the degree of

Doctor of Philosophy

in Economics

Acknowledgements

This thesis was written to a large extent thanks to the extensive support I received during my years as a doctoral student. First and foremost, I wish to extend my deepest gratitude to my dissertation advisor, Paul Seabright, without whom I probably couldn't have pursued a PhD in the first place. His vast knowledge in economics and beyond has been a limitless source of learning across the years, and my growth as an economist I owe mostly to our regular discussions. Fuelled by an exceptional scientific curiosity, Paul has kept constant interest in my research even when, at times, I lost the picture myself. I am grateful for this interest, as well as for his availability and encouragement.

Many other faculty members and colleagues had an important role in the writing of this thesis. This includes the participants to the BID workshop in Toulouse, who regularly offered valued feedback and constructive remarks. This also includes the organizers of and participants to the various ASREC events I attended, whose passion about shared academic interests have been an important motivator in my work. For all three chapters of this thesis, I benefitted greatly from stimulating discussions with François Salanié. Across our meetings I have come to recognize François's methodological skepticism and parsimony as formidable tools when tackling economic problems, a lesson which I will strive to carry with me in the future. I am also deeply thankful to Bernard Salanié, who hosted me at Columbia University during the spring 2022. Bernard's insightful advice on the first and most consequential chapter of this thesis was essential in shaping both my vision for this chapter and its realization. His time and his guidance are much appreciated. Lastly, I want to thank Audrey Richard-Ferroudji and Aurélie Varrel, who supervised my masters' thesis before I came to Toulouse, and who really instilled in me the taste for academic research.

In addition, I would like to thank the Institute for Advanced Study in Toulouse, the doctoral school of TSE, the TSE-R unit, and the Université Toulouse Capitole for

providing the necessary resources and environment to complete my research.

My friends and colleagues in Toulouse have been a great help in supporting my work, in more or less direct capacities. In this respect I wish to thank Alae Baha, Margot Dazey, Maxime Derex, Alipio Ferreira, Efisio Manunta, Paul-Henri Moisson, Romain Mongeau, Sébastien Montpetit, and Esteban Muñoz. I owe particular thanks to Vatsala Shreeti for encouraging and supporting me during the more difficult times, and for making my time at TSE and in Toulouse an unforgettable one.

Finally, I wish to thank my family, in particular my parents, for their unwavering support, patience, and understanding all along my studies.

Introduction

Within the broader context of societal norms, economic behaviors, and cultural practices, a question arises: how does one's cultural identity influence economic and social choices? This interplay between cultural imperatives and economic decisions is paramount in understanding individual motivations and societal trends, especially in a diverse and evolving landscape like France.

The urgency of this exploration is highlighted by France's unique challenges in navigating its rich cultural heritage while also accommodating the evolving identities within its borders. This thesis, delineated across three chapters, seeks to address specific aspects of these complex intersections.

The first chapter's focus on the transmission of religious beliefs, especially within minority groups, sheds light on the challenges and choices individuals face between cultural preservation and socio-economic progression. It forces a reconsideration of the economic implications of maintaining one's religious and cultural legacy.

The second chapter's investigation into veiling among Muslim women in France brings to the forefront the tangible economic consequences of cultural and religious practices. The veil illuminates the sacrifices and trade-offs that individuals make when cultural practices intersect with economic realities and societal perceptions.

Lastly, the intricacies of marital choices, when examined through market segmentation, offer a nuanced understanding of how societal norms and economic factors can jointly influence personal decisions. This chapter unveils the multifaceted factors that individuals weigh when making such intimate choices.

Below I provide more detailed summaries for each of these three chapters.

Chapter 1: Culture, human capital, and marital homogamy in France

What economic sacrifices are people willing to make to transmit their culture? Using data on religious affiliation in France, I study the intergenerational transmission of religion and how it interacts with children's educational outcomes. A reduced-form analysis suggests that mothers contribute to religious transmission more than fathers; religious minorities more than majorities; and lower-educated parents more than higher-educated ones. A mechanism that can explain these patterns is that higher-educated parents have a higher opportunity cost of transmitting their religion to their children. I investigate this mechanism through a structural model, in which parents endogenously decide their time investments in their child's culture on the one hand, and in their formal education on the other hand. The analysis suggests that heterogeneities in transmission patterns are driven primarily by heterogeneities in preferences for religious transmission across genders and religious groups, rather than by differences in parents' education. Furthermore, religious minorities pay a higher price for religious transmission in terms of their children's educational outcomes. For instance, by measuring this cost in terms of the probability that the child will obtain a college education, Muslim parents pay a cost between 8 and 13 times greater than that for Christians.

Chapter 2: Veiling and Economic Integration of Muslim Women in France (with Sébastien Montpetit)

The economic implications of policies limiting the wearing of the Islamic veil for Muslim women have largely been overlooked in many Western countries. This paper investigates the relationship between veiling behavior and economic participation using the largest sample of Muslim women in France. Firstly, we present new descriptive evidence about Muslim women in France, demonstrating a significant negative relationship between veiling and economic participation. Secondly, to disentangle the various motivations behind the joint decision to veil and to be economically active, we develop and estimate a discrete-choice model of veiling and labor force participation. Our findings indicate that veiled women are less economically active not only due to religious preferences but also because veiling substantially reduces their economic opportunities. Additionally, our results emphasize the significance of personal religious motives in the decision to veil, rather than community-based religious pressure. Consequently, our findings call into question the rhetoric used to justify policies that restrict the wearing of religious symbols in France.

Chapter 3: Are marriage markets segmented?

Are marriage markets segmented? In traditional matching models with transferable utility, marital assortativity is fully rationalized by differences in the surplus produced by each potential match. This model feature necessarily leads to overestimating the role of spouses' preferences in marital assortativity, at the expense of other explanations. In this paper, I study another mechanism behind marital assortativity which remains underexplored in the literature: the segmentation of marriage markets along spouses' traits. I extend the Choo–Siow model to account for market segmentation: individuals are assigned to submarkets with probabilities which depend on their gender and matching trait. Segmentation thus provides a new explanation for spousal assortativity, which is accompanied by a new division of the surplus between spouses compared to the Choo–Siow model. I study this framework using matching patterns on education in France. In order to identify the role of segmentation versus surplus in assortativity, I use exogenous variation in the segmentation of the marriage market provided by the termination of mandatory military service in France in 1996.

Collectively, these chapters provide a panoramic view of the confluence of culture, economics, and societal norms. By delving deep into these interrelationships, this thesis endeavors to shed light on the myriad ways in which individual choices and societal structures reciprocally shape one another in the vibrant tapestry of French society.

Contents

Chapter 1

Culture, human capital, and marital homogamy in France

1.1 Introduction

What sacrifices are people willing to make to transmit their culture? Whether implicitly or explicitly, cultural transmission shapes the trade-offs that people make on consumption and investment decisions on a daily basis.[1] Routine economic choices, such as how to dress their children, or whether the family attends a sports game rather than church on Sunday, are influenced by culture. But considerations of cultural transmission also influence major decisions in families' lives, such as which neighborhood to move to, or which school to enroll children in. Moreover, these considerations do not only apply to current parents. Even before having children, people anticipate how their choices will affect their ability to transmit their culture later on, in particular when they choose a partner. In turn, efforts to find a suitable partner might influence other life-changing decisions such as where to live, or whether to go to college. Thus, a wide range of critical choices and behaviors, which have independently been studied by economists for decades, are in fact shaped by cultural transmission. However, we still know very little about how and to what extent cultural transmission influences these economic decisions and outcomes.

In this paper, I address this issue by studying an important economic decision: parental investments in children's education; and its relationship with a crucial cultural trait:

[1] Intergenerational transmission is one of culture's defining features: following for instance Guiso, Sapienza and Zingales (2006), culture designates "those customary beliefs and values that ethnic, religious, and social groups transmit fairly unchanged from generation to generation."

religion. Specifically, I examine how parents trade off between intergenerational religious transmission and investments in their children's educational attainment, in the context of modern France.[2] The main argument of this paper is that parents from different religious and educational backgrounds face unequal trade-offs on this issue. In particular, religious minorities are more likely to invest in religious transmission at the expense of their children's educational attainment, and they pay a higher opportunity cost for it.[3] For instance, my results suggest that, at the margin, investments in religious transmission made by Muslim parents (the main religious minority in France) are between 8 and 13 times more costly than those made by Christian parents (the religious majority) in terms of the probability that their child will obtain a college degree.

To understand the nature of this trade-off, first I investigate the patterns of religious transmission and children's education by using French survey data from 2008. Both an extensive descriptive analysis and a reduced-form approach suggest that mothers invest in the transmission of their religious affiliation more than fathers, and that religious minorities (Muslims and Jews) invest more than majorities (Christians and Unaffiliated). Furthermore, lower-educated parents transmit their religious affiliation more successfully than higher-educated parents on average. Conversely, children of Christian parents are more educated than those of Muslim parents, even when controlling for the parents' education. The reduced-form analysis, which uses a multilogit specification to explain children's choice of religious affiliation as a function of their parents' characteristics and of the religious mix of their environment, successfully fits the data on parents' and children's religious affiliations.

In a second step, I explicitly address the trade-off between religious transmission and education by building a structural model in which parents must invest in the child's formal education, on the one hand, and in the child's religious socialization, on the other hand.

[2]The tension between religion and formal education has long been a particularly striking illustration of the trade-offs that cultural transmission entails. Modern schooling emphasizing rationality and the scientific approach has long clashed with religion, in part because their respective teachings are sometimes incompatible, but also because they must compete for children's limited attention (see for instance Squicciarini 2020, Chaudhary and Rubin 2011, or Carvalho, Koyama and Sacks 2017 for historical examples). In the United States this clash is still unfolding, for instance with the ever-lasting debates around the inclusion of creationism in the public school curriculum.

[3]Anecdotal evidence of this phenomenon is extensive. Some religious groups, such as the Amish or Jehovah's Witnesses, even explicitly discourage their affiliates from pursuing college or even high school education – arguably because these groups implicitly acknowledge these trade-offs. In September 2022, the *New York Times* reported on the dismal state of education in New York City's Hasidic Jewish schools, which have prioritized religious teachings at the expense of basic skills such as English and math ("In Hasidic Enclaves, Failing Private Schools Flush With Public Money", NYT, Sep. 11, 2022).

In this model, the trade-off arises because both the *socialization* process, whereby children learn the tenets and principles of the previous generation's culture, and the investments in the child's formal education, are time-consuming activities for the parents. Crucially, the model incorporates three key mechanisms which explain differences in how parents choose to invest in religion versus education for their children. The *first mechanism* is that higher-educated parents are more productive than lower-educated parents in furthering their children's formal education. This mechanism is directly inspired by the stylized facts derived from the reduced-form analysis, which suggest that higher-educated parents have a higher opportunity cost to transmit religion to their children. Here, I model this opportunity cost as foregone investments in the child's formal education. The *second mechanism*, called *cultural substitution*, is adapted from the literature on the economics of cultural transmission (Bisin and Verdier 2000). This mechanism entails that while parents from religious majorities can extensively rely on their environment to socialize their children, the same is not true for religious minorities. Consequently, religious minorities must invest comparatively more in religious socialization to achieve the same religious transmission outcomes. The *third mechanism* is preference heterogeneity across parents. I allow parental preferences for the child's religion versus education to vary across two dimensions, namely, parents' gender and religious affiliation. With these assumptions, I model parental behavior by using a collective household model, and I derive closed-form solutions for how parents invest in their children's religion versus education.

Finally, I estimate this model, leveraging the variation in children's religious affiliation and educational attainment. To exploit this double variation and to estimate parameters despite nonlinearities (for which standard logit regression is not suitable), I develop a maximum likelihood approach that combines elements from both multinomial logit (for religious affiliation) and ordered logit (for educational attainment) estimation. My results indicate that the three mechanisms discussed above matter for the parental trade-off between religious transmission and investments in their children's education, albeit at different scales. A log-likelihood decomposition analysis allows me to rank these three mechanisms by order of importance in terms of explanatory power. I find that parental preferences matter the most in explaining the variation in children's religious affiliation and educational attainment, followed by the economic mechanism involving a higher opportunity cost of religious socialization for higher-educated parents, and finally

by the cultural substitution mechanism.[4] Through counterfactual analysis, my estimation results also allow me to quantify the trade-offs that different parents face between investments in their child's religious socialization versus formal education. To do so, I use the estimates to reconstruct the households' production possibility frontier in terms of two household outputs, the religious transmission rate and the probability that the child will obtain a college degree. By measuring the slope of this frontier (i.e. the household's marginal rate of transformation) I recover the cost of religious transmission in terms of children's educational attainment, finding for instance that Muslim parents pay a cost 8 to 13 times greater than Christians parents.

These results have far-reaching implications for the way that we understand incentives, inequality, and education policy in relation to religion and, more broadly, to culture. First, they indicate that cultural minorities may have comparatively higher incentives to invest in cultural transmission for their children, over the acquisition of skills which are validated by diplomas and valued on the labor market. These incentives to invest in cultural transmission rather than education are likely to be reinforced by the fact that many cultural minorities typically face weaker job opportunities. The dynamic implications for inequality are severe, since these incentives would amplify any existing educational gap between cultural majorities and minorities across generations, on top of other structural reasons such as access to lower-quality public schools. Second, the fact that preferences play a large role in the trade-offs between culture and education is an important challenge for policy-makers. In this respect, an important policy objective is to conciliate formal education with cultural transmission for cultural minorities. There are many available options to advance this objective, such as public funding for denominational schools (accompanied by a proper amount of oversight on school curricula) and for cultural associations (which can take the burden of cultural transmission away from parents), or even revising the public school curriculum to make it more inclusive of pupils' diversity. My results suggest that such efforts could alleviate the educational gap between cultural minorities and majorities.

[4] Although I cannot rule out that parental preferences are in fact rooted in an economic value of children's religious affiliation, in the absence of further evidence it seems reasonable to interpret the results as religious transmission to children mattering *per se* to the parents. A reason for cultural affiliation to have an economic value could for instance be the existence of economic networks based on such affiliations; see Munshi (2011, 2019) on Indian caste-based networks. Starting with Iannaccone (1992), the economics of religion literature has also pointed out the 'club good' dimension of religion.

Contributions and related literature. By documenting how parents transmit their religion and human capital to their children in the context of contemporary France, this paper speaks to a recent literature that has explored investments in religious versus formal education in various settings and, more broadly, to the literature on the economics of religion (see Iannaccone 1998 and Iyer 2016 for reviews). For instance, Squicciarini (2020) shows that in 19th-century France, Catholic education competed with the secular curriculum in schools, ultimately hampering economic development in regions with higher religiosity. Chaudhary and Rubin (2011) and Saleh (2016) document a similar phenomenon for Muslims in colonial India and 20th-century Egypt, respectively. Carvalho, Koyama and Sacks (2017) and Carvalho, Koyama and Williams (2022) consider models in which cultural minorities protect their culture by resisting formal education, taking as illustration the 19th-century Jewish emancipation in Europe. Here, I contribute by exploring new reasons why parents may decide to invest in religion versus education, and by quantifying their effects.

As opposed to educational institutions, my paper focuses on how parents spend their time investing in religion versus education for their children, and in this respect it fits within the literature on time allocation theory and the human capital formation of children (Becker 1965, Cunha and Heckman 2007). Indeed, parental time investments have been shown to be important factors in children's human capital formation (Del Bono et al. 2016) and cultural capital formation (Botticini and Eckstein 2007, 2012, Patacchini and Zenou 2016). In particular, my model of cultural socialization takes inspiration from the technology of children's human capital formation in Del Boca, Flinn and Wiswall (2014, 2016). I use a collective household framework (Chiappori 1992) to model parental time investment decisions. In that respect, a paper close to mine is Chiappori, Salanié and Weiss (2017), which models trade-offs between time investments in children's human capital and time spent working in order to explain the evolution of the marital college premium. Here, I aim to explain heterogeneities in the patterns of religious and human capital transmission, and to that end I adapt this model to focus instead on time investments in children's religious versus human capital formation. For this purpose I consider religious capital as an intensive measure of religion[5] which is built by purposeful investments. This approach

[5]Rather than just religious affiliation, which is a discrete, extensive measure. The literature on cultural transmission has mostly focused on the latter for now, but see Cheung and Wu (2018) who consider a continuous trait and Patacchini and Zenou (2016) who consider discrete intensity (low or high religiosity) of the same trait. This focus on intensity has remained at the expense of a possible multiplicity of traits,

can be traced to Iannaccone (1990) who initially considered a human capital approach to religion.

Finally, this paper relates most directly to the literature on the economics of cultural transmission spurred by Bisin and Verdier (2000, 2001) and reviewed twice since then (Bisin and Verdier 2011, 2022). On the empirical side, it joins other works that use cross-sectional data on parental and children's cultural affiliations to recover values for the primitive parameters of cultural transmission models.[6] Close papers include Bisin, Topa and Verdier (2004), who also study the transmission of religious affiliation but using US data; Patacchini and Zenou (2016) who study parental religious socialization efforts as a function of the child's religious environment, also in the US; or Bisin and Tura (2022) who study language transmission among Italian migrants. Methodologically however, these papers rely mainly on aggregate moments (probability of transmission or of homogamous marriage) to estimate structural parameters. As far as I know, using discrete choice theory (McFadden 1973) to empirically explain children's choice of cultural affiliation, as I do in this paper, is a new contribution to this strand of literature. This methodological shift reflects the fact that I also depart from the usual Bisin and Verdier framework, which focuses on discrete cultural affiliations, for an approach that emphasizes cultural capital formation. Another recent effort to include a cultural capital approach in the theory of cultural transmission is Carvalho and McBride (2022). In their model, parental socialization investments contribute to determining the child's cultural type (extensive margin). Later, children can then build their cultural capital upon this type (intensive margin). This differs from my model, wherein parental investments directly contribute to the cultural capital, from which individuals derive their type. Furthermore, starting with Bisin and Verdier (2000) the cultural transmission literature has mostly considered costs to cultural socialization efforts in an abstract way. Here I contribute by considering a very concrete cost, namely, the time opportunity cost of socialization on investments in children's education. This allows me to measure the cost of religious transmission in terms of children's educational attainment, an economic outcome of primary concern. Finally, by using a collective household model to explain socialization decisions I depart from the standard unitary model used in the cultural transmission literature (an approach also

however.

[6] In their review, Bisin and Verdier (2011) label these papers collectively as *structural socialization studies*.

taken recently by Bisin and Tura 2022). Most notably, this modelling choice lets me model the behavior of heterogamous households in a non-trivial way, and identify the separate contributions and characteristics of mothers and fathers in the transmission process – something which is not possible with the unitary model.

The paper is organized as follows. In section 2.3 I describe the data along several dimensions of interest: education, religion, and patterns of marriage and of intergenerational transmission. In section 1.3 I use reduced-form analysis to focus on empirical patterns of religious socialization. In section 3.2 I introduce the theoretical framework of cultural socialization in the household, in which parents must trade off investments in their child's culture versus formal education. In section 1.5 I describe my procedure for estimating this model, and present my results. Finally, section 3.6 concludes.

1.2 Data

To investigate the relationship between culture and human capital in marriage and transmission to children, I use data from the Trajectories and Origins survey (*Trajectoires et Origines*, or TeO for short; see Beauchemin et al. 2016 for details). The TeO survey was conducted in metropolitan France in 2008. With over 21,000 respondents, it aimed to document the life experiences of migrants living in France and their descendants. Because of this specific aim, the TeO survey is particularly relevant for studying intergenerational transmission. First, it includes questions not only about the respondents, but also about their parents. This information is obviously critical to the study of intergenerational transmission. Second, it is one of the few large-scale surveys in France that collects answers on respondents' religious affiliation and practices. Indeed, collecting such information is generally prohibited by law in France (*loi informatique et libertés* of 1978) and requires a special derogation. For the purpose of this paper, it means that the TeO database is a rare opportunity to study religion as an example of cultural trait in France. Last, the TeO survey oversamples migrants and their descendants by design. In doing so, it provides a sizeable sample for several religious minorities in France, most notably Muslims, thus allowing me to draw comparisons across different religious groups.

Respondents were between 18 and 60 years old at the time of the survey (cohorts born between 1948 and 1990). The sample is slightly skewed toward women (52.8%). In the

following, not only do I use data on the respondents themselves, I also rely extensively on the answers regarding their parents to study time trends, as well as marriage and transmission patterns. Respondents' parents were born as early as 1900, but I ignore pre-1920 parental cohorts on graphical representations (those have fewer than 100 observations per cohort). I provide more general statistics about the TeO survey in Table 1.1 (Appendix 1.7.1). In the rest of this section, I describe the TeO data and some stylized facts regarding education and religion, in terms of both transmission and marriage patterns.

1.2.1 Education

In the TeO survey, educational attainment is reported through the International Standard Classification of Education (ISCED) 1997. From this variable, I construct three simplified educational attainment categories: (1) "Primary or less," for individuals who completed at most primary education; (2) "Secondary," for individuals who obtained a middle- or high-school diploma, or a technical diploma from an age-equivalent training program; (3) "Tertiary or more," for individuals who hold a postsecondary diploma. The proportions of these categories in the respondent sample are 8% (primary or less), 64% (secondary), and 28% (tertiary or more). Among the respondents' parents, these proportions are 57%, 31%, and 12%, respectively.

Educational attainment. Figure 1.1 shows the evolution of educational attainment by gender for the 1920–1978 cohorts, mixing data on respondents and their parents.[7] In this figure I omit the youngest cohorts, who may not yet have completed their education at the time of the survey. (I chose the 1978 cohort, who was 30 years old at the time of the survey, as the endpoint.) Educational attainment increases for both genders across the cohorts under study. Beginning approximately with the 1970 cohort, women overtake men in tertiary education.

[7] A note on graphical representations: the data can be quite noisy and the graphics difficult to interpret when observations are split across the three dimensions of cohorts, religion, and education. For this reason, graphical representations throughout the paper feature nonparametric predictions (LOESS) of different outcomes of interest on birth cohorts. This approach allows me to obtain smoothed curves that provide a better picture of the evolution of these outcomes across cohorts (see for instance Figure 1.1 on education in the sample). These curves are systematically accompanied by representations of the corresponding 95% confidence interval. On some graphs I represent the actual data with dots (such as in Figure 1.1), but when doing so would hamper readability I represent only the nonparametric predictions (such as in Figure 1.4).

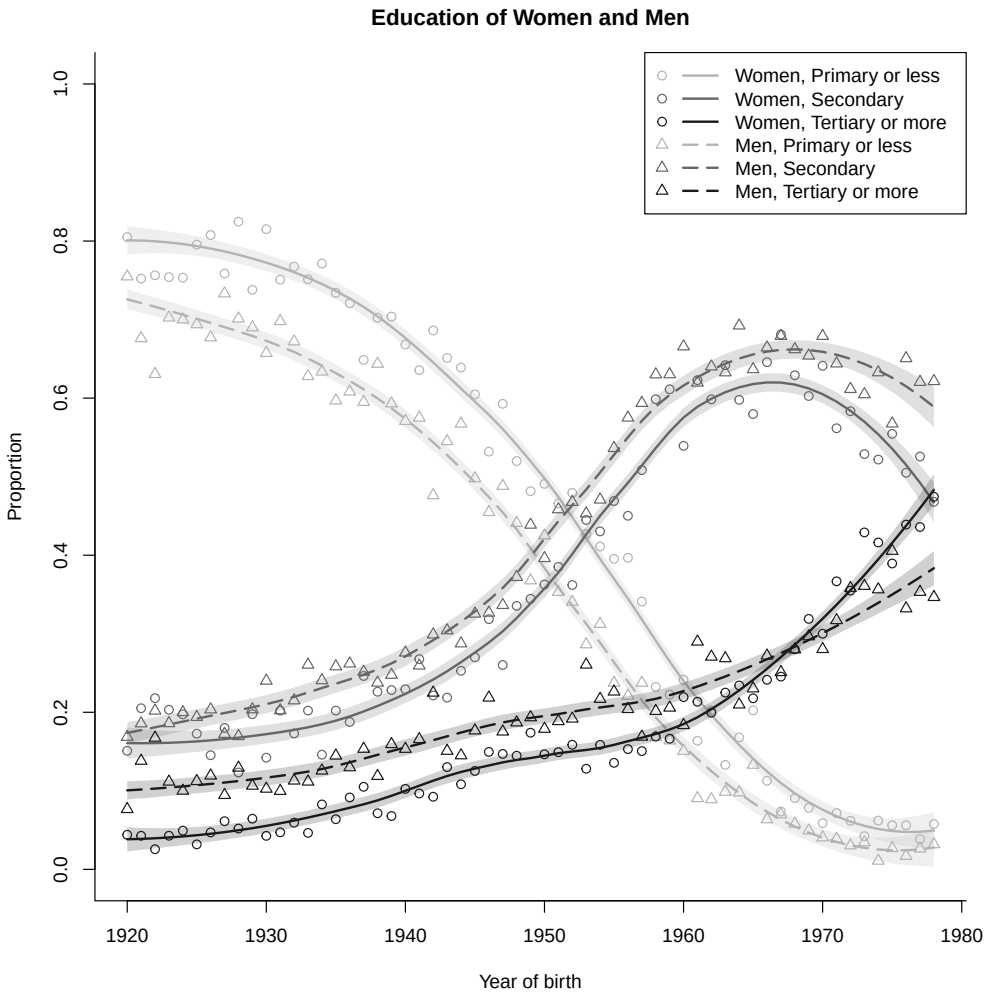

Figure 1.1: Education of women and men.

Marital assortment. Although 72.5% of respondents declared that they had a partner, information on this partner was collected only when they lived in the same house (60.9% of respondents). Once again, I also use answers on the respondents' parents to draw a long-term picture of marital assortment in the sample, which I present in Figure 1.2. We can discern some time trends in educational homogamy. The proportion of couples with the same educational attainment is high overall. It decreases for the oldest cohorts, from 80% in 1920 to approximately 65% in 1950. This decrease might be simply a mechanical consequence of the increasing diversification of educational attainments for these cohorts (early cohorts mostly had only a primary education, so there could not be many mixed-education couples). After 1950, this proportion stagnates between 65% and 70%. The proportion of couples with a more educated husband increases slightly across the oldest cohorts, and then starts to decrease around the 1950 cohort to reach 15%. The proportion

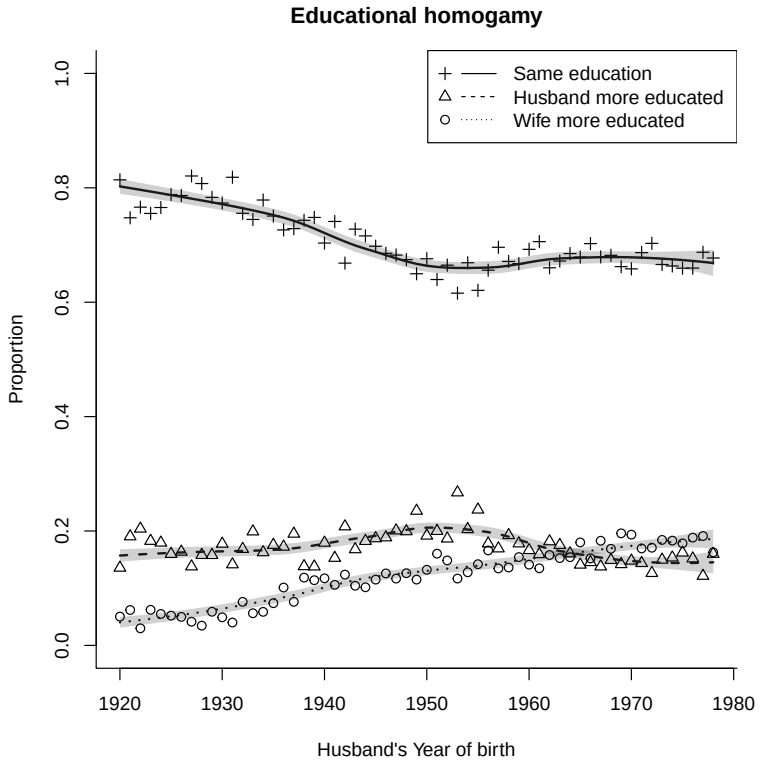

Figure 1.2: Educational homogamy.

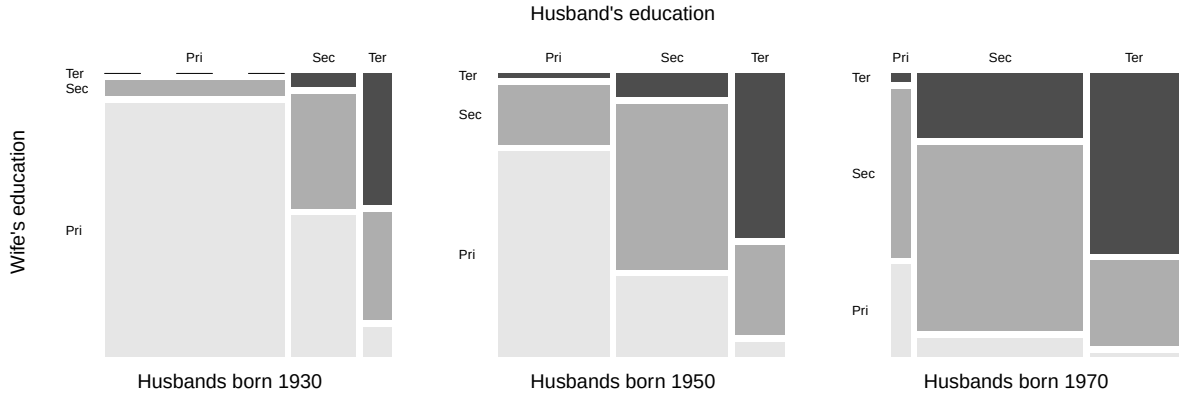

Figure 1.3: Educational assortment in couples with a husband born in 1930, 1950, and 1970.

of couples with a more educated wife increases across all cohorts, from 5% to almost 20%, overtaking the proportion of couples with a more educated husband by the 1965 cohort.

Could these trends be driven by the simplification of the education variable into three categories? In Figure 1.1 I construct the same graph with the detailed diploma categories (8 levels, from no diploma to university graduate). While the proportion of couples with the same educational attainment mechanically falls when considering more education levels, the trends discussed above mostly hold. In particular, the proportion of couples

with a more educated wife clearly increases over the cohorts considered, overtaking the proportion of couples with a more educated husband.

Finally, in Figure 1.3 I report detailed educational assortment patterns for three different cohorts, defined as those couples with a husband born in 1930, 1950, or 1970. (Figure 1.3 does the same for wives born in 1930, 1950, or 1970.) In accordance with the trends discussed above, we observe that more men marry "up" among younger cohorts (the number of educated women has increased more than the number of educated men). As Figure 1.2 already suggested, by the 1970 cohort marriage patterns are almost symmetric for men and women: approximately as many women marry up as men do.

In Appendix 1.7.5 I also study educational homogamy through local log-odds ratios, following Siow (2015). Statistical tests of the TP2 criterion (i.e. Total Positivity of order 2 for the local log-odds ratios) provide strong evidence of educational homogamy, both on the complete sample and conditional on spousal religious affiliations.

1.2.2 Religion

In the survey, religious affiliation is recorded via 13 possible answers. To simplify the analysis, and because some answers are associated with few observations, I aggregate them into five broad categories: No religion or "Nones" (29% of respondents), Christian (39%), Muslim (27%), Jewish (1%), and Other religion (4%).

Religious affiliation. Figure 1.4 presents religious affiliation across cohorts in the sample by gender (including parents). For both genders, younger cohorts show a higher representation of Muslims and Nones, and fewer Christians. The representation of Jews and other religious affiliations remains low and stable across cohorts. Interestingly, while Muslims, Jews, and Others are well balanced in terms of gender, there is an important asymmetry among Christians and Nones. Indeed, Christian women are more numerous than Christian men (48% of women vs. 43% of men are Christian). This difference is almost perfectly balanced by the excess of None men compared to None women (24% vs. 19%).

It is worth noting that religious affiliation in the sample is not especially representative of the French religious mix (for instance, 27% of respondents identified as Muslim, even though usual estimates for the share of Muslims in France hover between 5% and 10%

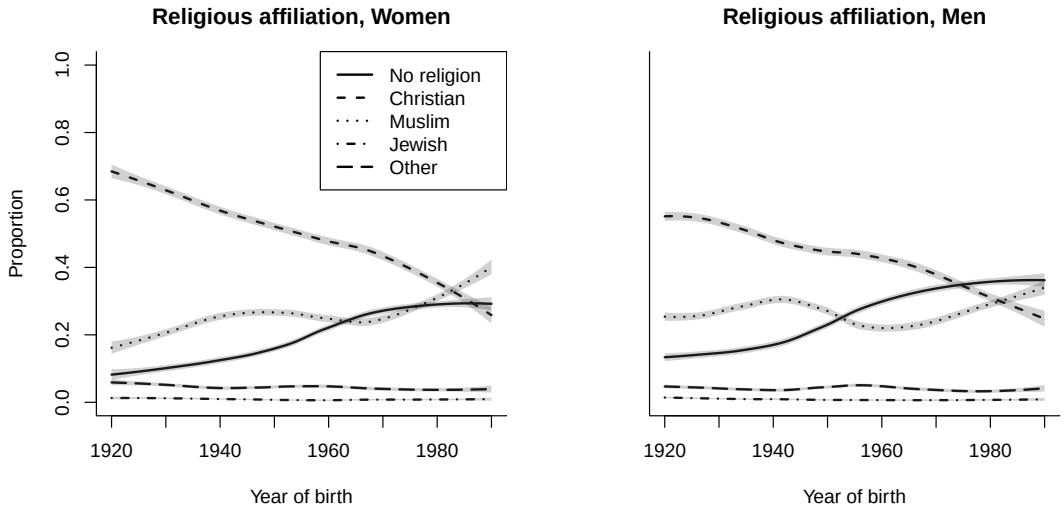

Figure 1.4: Religious affiliation, women and men.

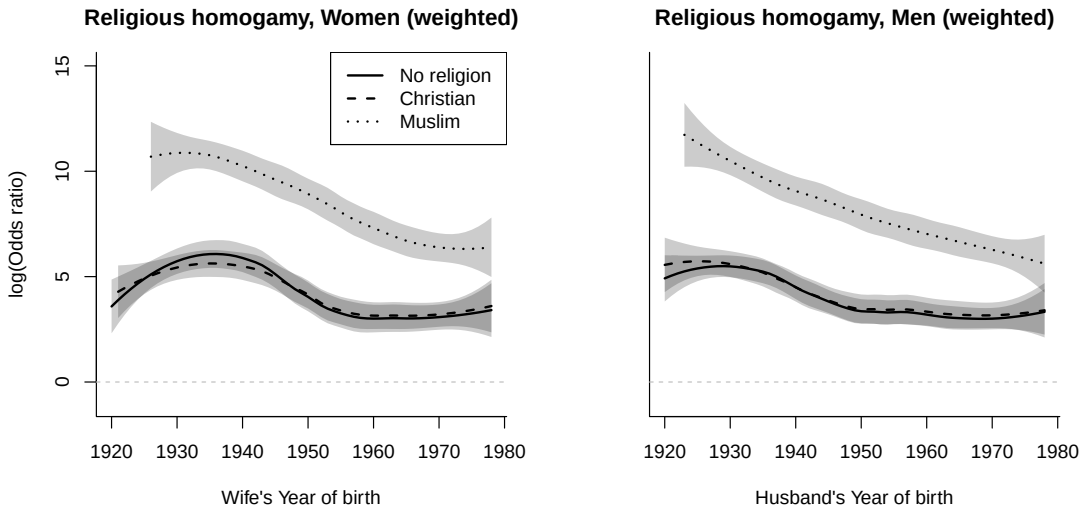

Figure 1.5: Religious homogamy (log-odds ratios), Women and Men.

for 2008). This bias is a natural consequence of the TeO survey oversampling individuals with a family history of immigration. Figure 1.5 reproduces the graphs of Figure 1.4 by using the sampling weights provided by the survey, providing a better (but still imperfect) picture of the share of each religious affiliation in France.

Marital assortment. A common way to measure partner assortativity along one dimension (here, religious affiliation) is to compute the log-odds ratios:

$$\ln\left(\frac{n_{aa}\, n_{\bar{a}\bar{a}}}{n_{a\bar{a}}\, n_{\bar{a}a}}\right),$$

where n_{aa} is the number of individuals from affiliation a with a partner a, $n_{a\bar{a}}$ that of individuals a with a partner non-a, and so on. Log-odds ratios are equal to 0 when couples are formed randomly,[8] while positive log-odds ratios are evidence of homogamy (positive assortative matching), and negative log-odds ratios are evidence of heterogamy (negative assortative matching).

Figure 1.5 presents these log-odds ratios for any birth cohort of husbands and wives, considering sampling weights. All computed log-odds ratios are positive for the cohorts considered, providing evidence of strong religious homogamy in the sample. Assortativity is stronger among Muslims than among Christians or Nones, although it decreases across cohorts: younger Muslims are less prone to religious homogamy than older Muslims. Christians and Nones exhibit similar and stable rates of homogamy from the 1950 cohort onward. (There is a decline in homogamy rates for these affiliations from approximately 1950, but this decline could be due to selection issues with the parents' generation in the sample.) I have omitted the log-odds ratios for Jewish and Other religions, since these affiliations have few observations per cohort, resulting in noisy patterns. It is however worth noting that despite this noise, both these affiliations exhibit high average assortativity rates that are closer to those for Muslims than for Christians or Nones. Figures 1.6 and 1.7 show assortativity patterns for the three cohorts born in 1930, 1950, and 1970, and provide further evidence of strong religious homogamy in the sample. Table 1.2 presents the 2×2 matrix of couples by religious affiliation.

Education by religious affiliation. Educational attainment is not distributed equally among religions, as shown in Figure 1.6. While religious Nones and Christians exhibit similar levels of educational attainment for the cohorts considered, Muslims have lower ed-

[8] When couples are formed randomly, individuals a and non-a have the same odds of being matched with a partner a over a partner non-a, i.e.

$$\frac{n_{aa}}{n_{a\bar{a}}} = \frac{n_{\bar{a}a}}{n_{\bar{a}\bar{a}}}.$$

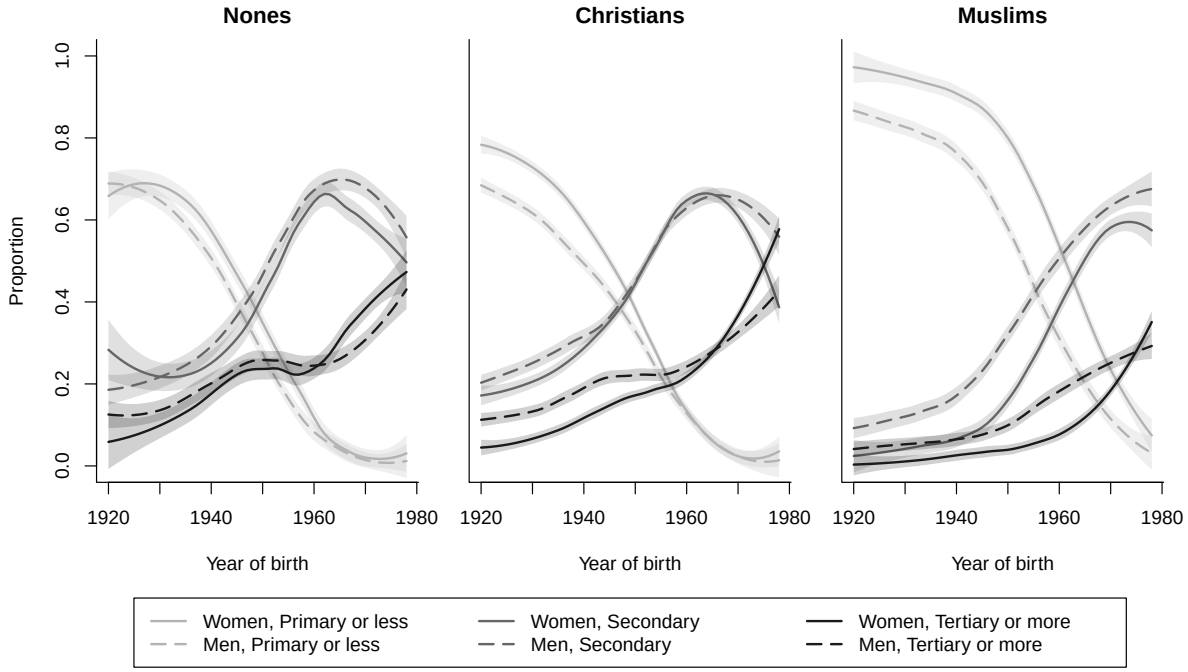

Figure 1.6: Education by religion and gender.

ucational attainment throughout. Regarding the interaction between gender and religion, for the oldest cohorts (1920–1950) men are more educated across all religions. Beginning with approximately the 1950 cohort, this gender gap starts to close among Christians and Nones (a slight gender gap in favor of women even appears among Nones), while it persists until 1970 among Muslims. It is only for the very latest cohorts that a discernable gender gap appears in favor of women for all three religious affiliations.

1.2.3 Marital assortment on education and religion

On education conditional on religion. Figure 1.7 presents the patterns of educational assortment for same-religion couples. Nones and Christians exhibit similar patterns of high educational assortment: partners have the same education level in approximately 70% of couples, although this rate decreases slightly over the cohorts considered. Muslims show a greater proportion of couples for which the husband is more educated, but this is expected as a mechanical consequence of the educational gap in favor of men in that population, as discussed above. Again, as in the case of Figure 1.2 for the complete sample, I verify that these results hold when considering more detailed diploma categories (see Figure 1.2).

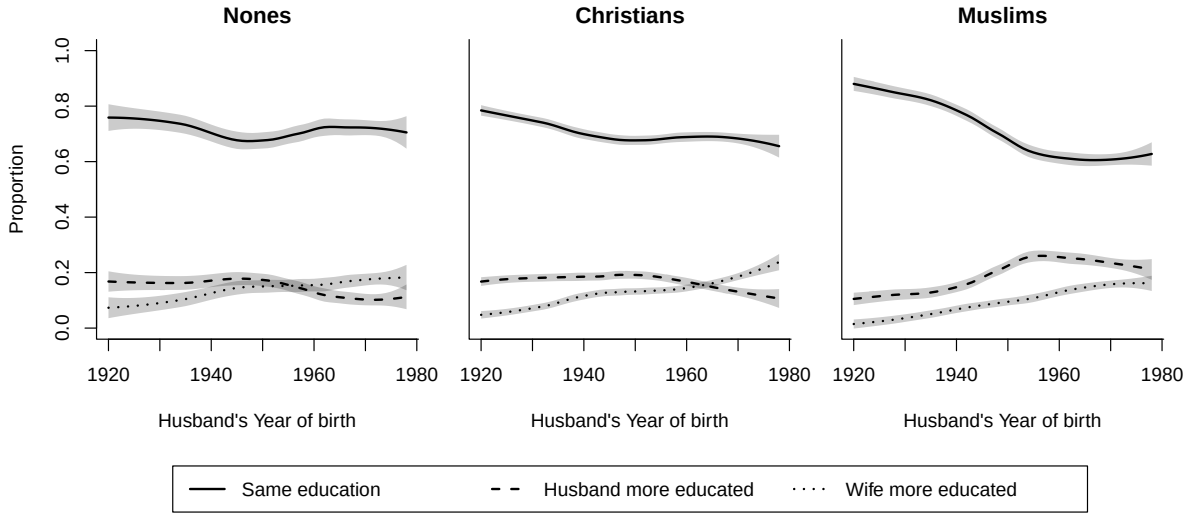

Figure 1.7: Educational homogamy, same-religion couples.

On religion conditional on education. Figure 1.8 again shows log-odds ratios, but this time compares education levels to see how they might affect religious assortment. In the complete sample (first row of the graph), there appears to be a negative correlation between religious homogamy and educational attainment: religious homogamy is strongest among individuals with a "Primary or less" education, and weakest among individuals with a "Tertiary or more" education. This difference could however be due to a correlation between religious affiliation and educational attainment (we have seen for instance that Muslims in the sample are simultaneously less educated and more homogamous on average). To alleviate this concern, I examine how religious homogamy differs across education levels within religious affiliations (rows 2 to 4 of the graph). Importantly, the evidence becomes fragmentary when considering such interactions, because data become thinly spread across categories. However, the negative correlation between religious homogamy and educational attainment seems to hold within religious affiliations. It is most pronounced for Muslims, as well as for None women and Christian men, and especially for older cohorts. The correlation is less pronounced for Christian women and None men, who have overall lower levels of homogamy because of the asymmetry documented above (there are more Christian women than men, and more None men than women).

Table 1.1 confirms this negative correlation between religious homogamy and partners' education with a simple linear regression. According to the second specification, which includes religion fixed effects for both partners, an increase in the husband's or wife's

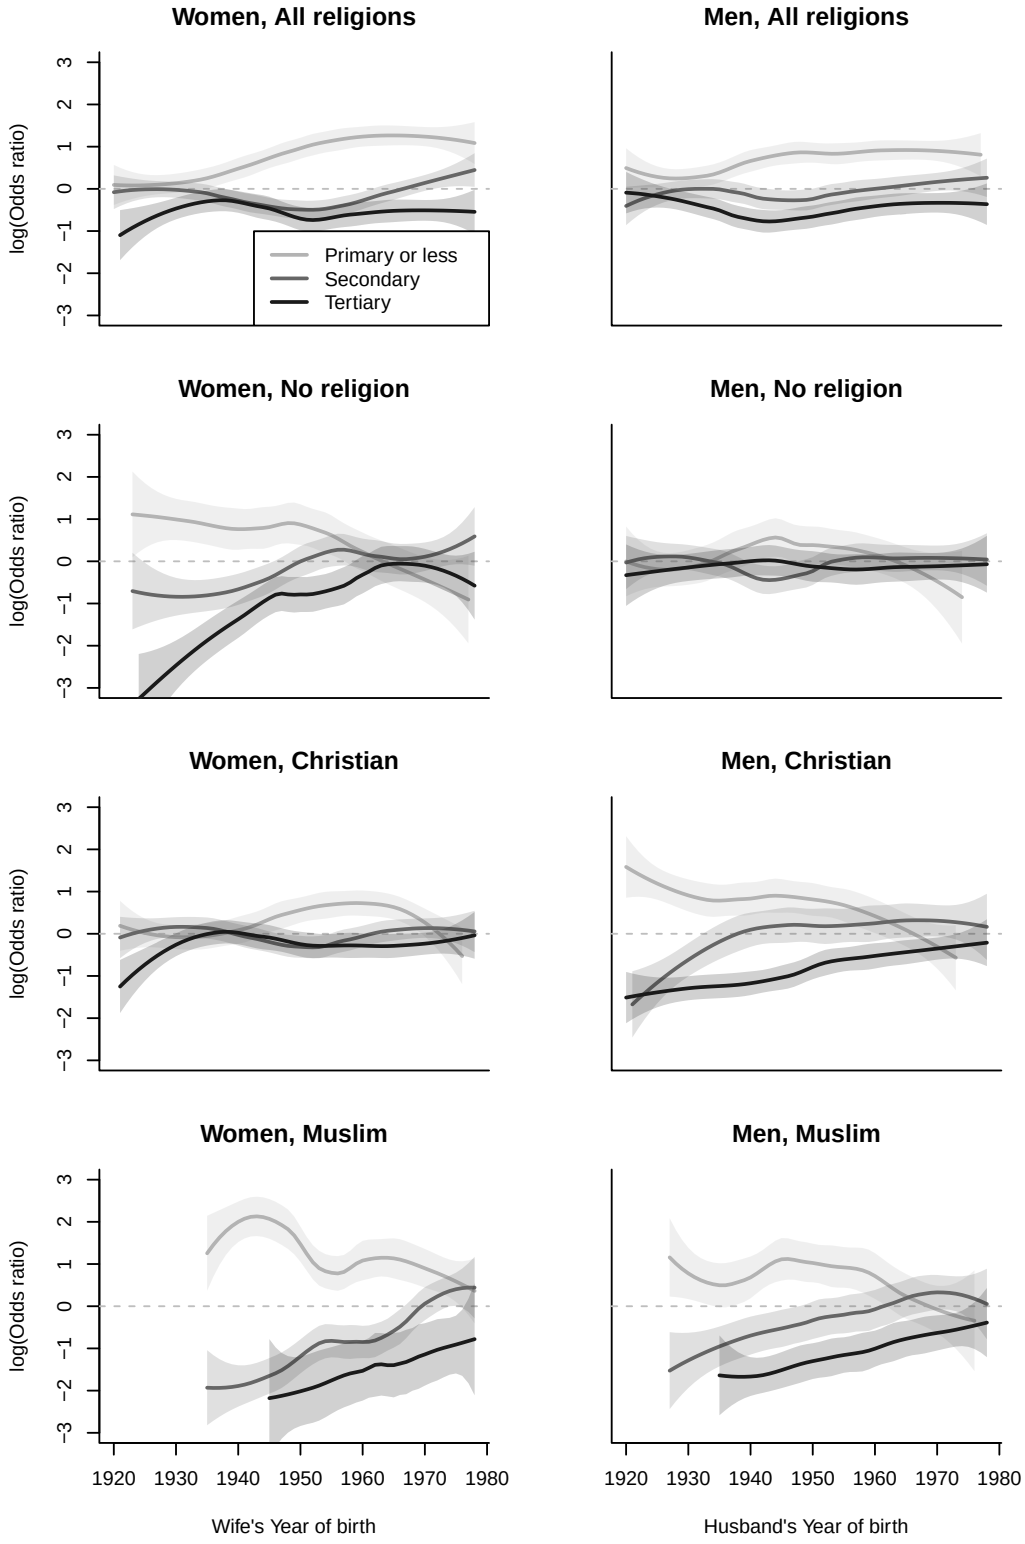

Figure 1.8: Religious homogamy of women and men, by Religion and Education. Here, log-odds ratios are computed within a religion category. For instance, in the "Women, Muslim" graph the "Primary or less" line is obtained by computing the odds of a Primary–Muslim woman being partnered with a Muslim man, divided by the odds of a Secondary– or Tertiary–Muslim woman being partnered with a Muslim man.

Table 1.1: Religious homogamy and education.

	Religious homogamy	
	(OLS)	(OLS)
Wife's education		
Secondary	−0.05*** (0.01)	−0.03*** (0.01)
Tertiary	−0.06*** (0.01)	−0.04*** (0.01)
Husband's education		
Secondary	−0.04*** (0.01)	−0.03*** (0.01)
Tertiary	−0.04*** (0.01)	−0.04*** (0.01)
Wife's religion		
Christian		−0.30*** (0.01)
Muslim		0.16*** (0.02)
Jewish		0.15*** (0.04)
Other		−0.38*** (0.03)
Husband's religion		
Christian		0.40*** (0.01)
Muslim		−0.06*** (0.02)
Jewish		−0.26*** (0.03)
Other		0.14*** (0.03)
Observations	31150	31150
Sampling weights	Yes	Yes
Adjusted R^2	0.01	0.17

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Reference category for wife/husband education fixed effects is "Primary."
Reference category for wife/husband religion fixed effects is "No religion."

educational attainment from Primary to Secondary is associated with a 3 p.p. decrease in the probability that he or she belongs to a homogamous couple. An increase from Secondary to Tertiary leads to a 1 p.p. decrease in the same probability.

1.2.4 Transmission of education

Education of the parents. Unsurprisingly, the children of higher-educated parents have higher education themselves, as seen in Figure 1.9. This finding is confirmed in Table 1.2 by an ordered logit regression with the child's education as the outcome, and the parents' educational attainment as the main explanatory variable. Corresponding specifications that use a linear model instead of an ordered logit model yield similar results (see Table 1.4).

Religion of the parents. To see whether religion plays a role in the transmission of education, in Figure 1.10 I plot the educational attainment of children as a function of

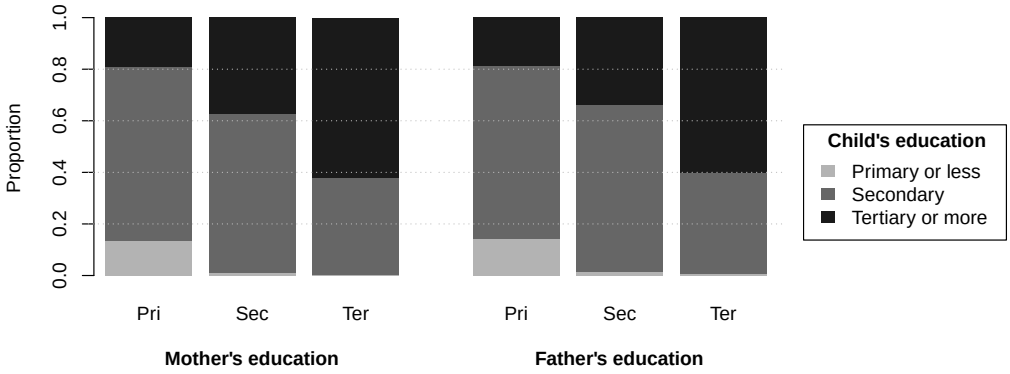

Figure 1.9: Transmission of education.

Table 1.2: Transmission of education (Ordered logit).

	Child's education		
	(Ord. logit)	(Ord. logit)	(Ord. logit)
Mother's education			
Secondary	0.64*** (0.02)	1.04*** (0.03)	0.97*** (0.03)
Tertiary	1.00*** (0.03)	2.60*** (0.10)	2.57*** (0.10)
Father's education			
Secondary	0.63*** (0.02)	0.99*** (0.03)	0.96*** (0.03)
Tertiary	1.56*** (0.03)	1.74*** (0.05)	1.72*** (0.05)
Mother's × Father's education			
Secondary × Secondary		−0.76*** (0.04)	−0.68*** (0.04)
Secondary × Tertiary		−0.40*** (0.06)	−0.36*** (0.07)
Tertiary × Secondary		−2.03*** (0.11)	−2.10*** (0.11)
Tertiary × Tertiary		−1.76*** (0.12)	−1.75*** (0.12)
Mother's religion			
Christian			0.25*** (0.02)
Muslim			0.03 (0.28)
Jewish			−0.02 (0.16)
Other			0.63*** (0.16)
Father's religion			
Christian			0.28*** (0.02)
Muslim			−0.11 (0.28)
Jewish			1.23*** (0.16)
Other			−0.74*** (0.21)
Child's year of birth /100	0.30*** (0.06)	0.40*** (0.06)	1.07*** (0.07)
Cut-off: Primary → Secondary	3.56 (1.24)	5.58 (1.25)	19.03 (1.28)
Cut-off: Secondary → Tertiary	7.70 (1.24)	9.76 (1.25)	23.25 (1.28)
Observations	18 793	18 793	18 222
Sampling weights	Yes	Yes	Yes
Deviance ($-2 \ln L$)	27 098	26 947	25 901

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Reference category for mother/father education is "Primary."

Reference category for mother/father religion is "No religion."

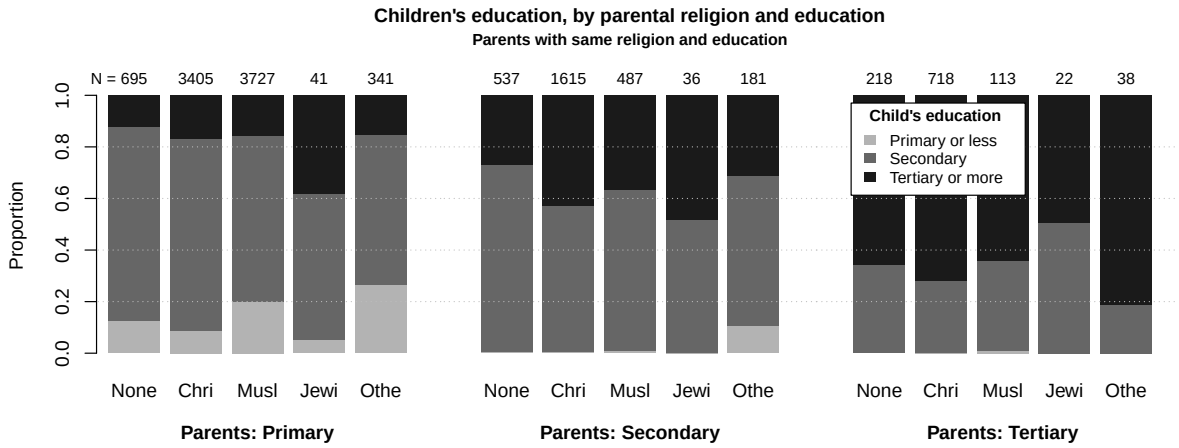

Figure 1.10: Transmission of education by parental characteristics.

both parental religion and parental education (focusing on parents who have the same religion and education). This approach shows that even if we hold the education of the parents fixed, the children of Christian parents tend to be more educated on average, as do the children of Jewish parents (although that sample size is much smaller). In contrast, children of None parents and Muslim parents seem to have lower education.

To inquire further, I include parents' religion as an explanatory variable in the previous ordered logit regressions of Table 1.2. The religious affiliations of the parents do seem to play a role in the transmission of education. Compared to the "No religion" baseline, Christian and Other mothers, and Christian and Jewish fathers, are associated with higher-educated children. Conversely, Other fathers are associated with lower-educated children. Note that these results might be dependent on patterns of religious homogamy. For this reason, I add interactions between mothers' and fathers' religious affiliations as explanatory variables (due to the lack of space given the numerous interactions, I only report these results in the Appendix, Table 1.3). The estimates of Table 1.2 remain robust when these interactions are added, while the estimated coefficients for these interactions are generally not statistically significant. At this point, it is however impossible to say whether these potential differences in children's education across parental religious affiliations are driven by trade-offs between religious socialization and education, or (for instance) by different cultural preferences for children's education. The structural model will address this question in section 3.2.

1.2.5 Transmission of religion

Homogamous vs. heterogamous couples. It is well-documented that parents in homogamous couples (i.e. couples in which both parents have the same religious affiliation) pass on their religion more reliably than parents in heterogamous couples (see e.g. Bisin and Verdier 2000, p. 960). This stylized fact remains true for the TeO data, as suggested by Figure 1.11. The *transmission rate*, defined as the probability that a child will have the same religion as their parent, is more than 80% among homogamous couples, and increases slightly across cohorts. This increase could simply be due to the change in the religious mix of the sample, with more Nones, more Muslims, and fewer Christians among younger cohorts (cf. Figure 1.4). However, other explanations are possible: younger cohorts could transmit more accurately, or this increase could be the result of individuals switching affiliation during their lifetime, so that older individuals would be less likely to continue to share their parents' affiliation. In contrast, the transmission rate for mothers and fathers in heterogamous couples is approximately 40%, half that of homogamous couples.

Figures 1.10 and 1.11 also describe religious transmission patterns across parental religious affiliations, this time aggregating all cohorts. They provide evidence for a homogamy advantage in religious transmission, except for the None affiliation.

By religion. Figure 1.12 presents transmission rates of mothers and fathers of the three main religious affiliations. Muslim transmission rates are higher overall, which might in part be a consequence of stronger homogamy among Muslims. However, Muslim transmission rates have also increased across recent cohorts, despite decreasing Muslim homogamy rates (cf. Figure 1.5). There are at least two possible explanations for this increase. First, since the population share of Muslims has increased over the period, it is possible that oblique socialization has become a better vector of religious transmission for Muslims. Second, as already discussed above, older individuals may be more likely to have switched affiliation from the one they inherited from their parents. The transmission rates of None parents follow a similar pattern, and are subject to the same interpretations.

In contrast, Christian transmission rates are decreasing, falling behind Muslim and None transmission rates beginning with the 1960 cohort. Two facts discussed above might contribute to this decrease: the population share of Christians is decreasing (cf. Figure 1.4), which may worsen oblique socialization, and homogamy rates among Christians are

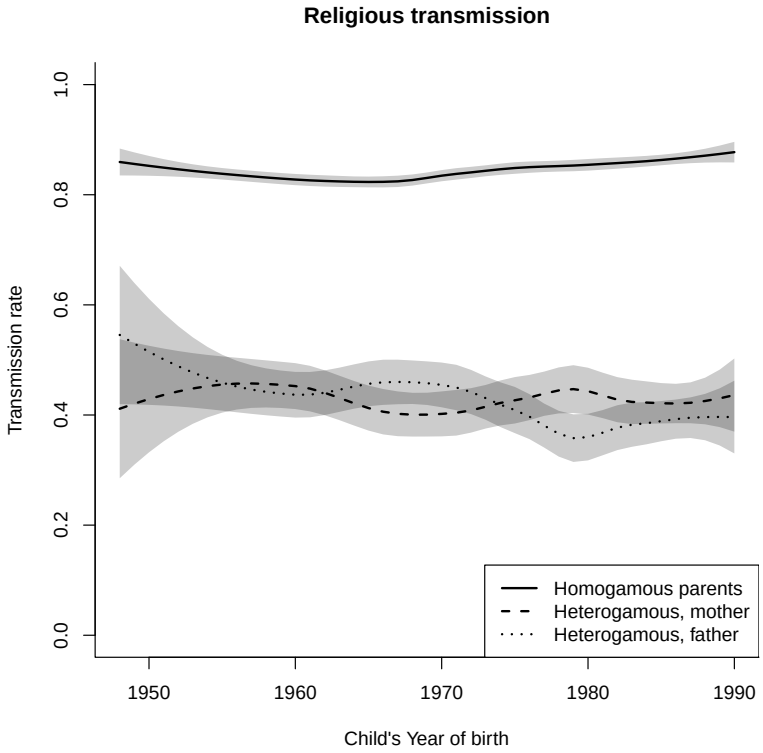

Figure 1.11: Religious transmission by parents in homogamous and heterogamous couples.

decreasing for the parental cohorts.

There are also comparisons to draw between mothers' and fathers' transmission rates. First, Christian mothers have lower transmission success than Christian fathers. A possible explanation for this difference is the asymmetry in the religious distribution of men and women. Indeed, since there is an excess of Christian women compared to Christian men, more Christian women end up partnered in heterogamous couples (most often, with None men), thus hurting their transmission rate. Conversely, None mothers have higher transmission success than None fathers, for the opposite reason: there is an excess of None men compared to None women. For Muslims, for whom there is less distributional gender asymmetry, there is no such stark difference between mothers' and fathers' transmission rates.

By educational attainment. Does the education of the parents matter in the transmission of religious affiliation? Figure 1.13 shows the transmission rates of mothers and fathers by educational attainment for all religious affiliations combined, and then separately for Nones, Christians, and Muslims. Despite the noise (data become thinly spread

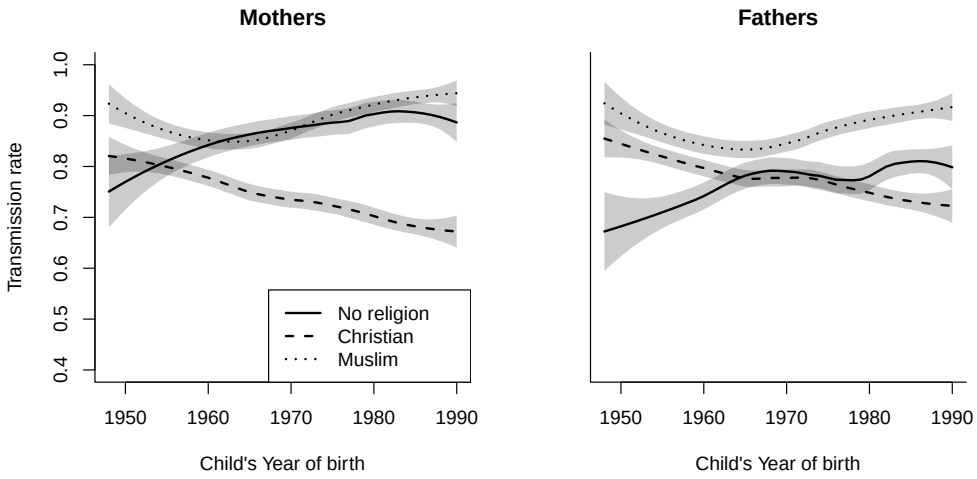

Figure 1.12: Religious transmission by mothers and fathers.

across the four dimensions considered: gender, birth cohort, religion, and education), the pattern that emerges is that parents with lower educational attainment have higher transmission rates. This finding is relatively clear when all religions are combined. When considering specific affiliations, the educational gap in the transmission rate is most pronounced for Muslims, and least pronounced for Nones. This result closely mirrors the pattern observed for homogamy rates and partner education (cf. Figure 1.8). For this reason, from Figure 1.13 it is unclear whether education affects transmission rates directly, or through its effect on religious homogamy. We can alleviate this concern by restricting attention to homogamous households only. In Figure 1.14, I present the transmission rates for mothers and fathers in homogamous households, excluding Nones. The pattern observed above persists: transmission rates are negatively correlated with parental education.

To clarify this finding I perform a simple linear regression and report the results in Table 1.3. Fathers' educational attainment is negatively correlated with the transmission rate (thus conforming to the pattern observed in Figure 1.13), consistent with the finding on homogamy: higher-educated fathers marry less homogamously, and thus can be expected to transmit religion less accurately. In contrast, mothers' educational attainment is positively correlated with the transmission rate. This positive correlation might seem puzzling: higher-educated mothers marry less homogamously and yet transmit religion more accurately. Note also that parents' education by itself has very little explanatory power for the transmission rates, as measured by the adjusted R^2.

Table 1.3: Religious transmission and education of the parents.

	Transmission rate, Mother			Transmission rate, Father		
	(OLS)	(OLS)	(OLS)	(OLS)	(OLS)	(OLS)
Mother's education						
Secondary	0.00	0.00	0.00	0.00	0.00	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Tertiary	0.04***	0.04***	0.05***	0.02	0.04***	0.06***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Father's education						
Secondary	−0.03***	−0.02**	−0.01	−0.02***	−0.01**	−0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Tertiary	−0.07***	−0.05***	−0.05***	−0.06***	−0.05***	−0.04***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Child's year of birth /100			−0.13***			−0.11***
			(0.03)			(0.03)
Mo.'s × Fa.'s religion FE		✓	✓		✓	✓
Observations	18 343	18 115	18 115	18 175	18 115	18 115
Sampling weights	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.00	0.12	0.12	0.00	0.11	0.11

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Reference category for mother/father education is "Primary."

A first step to disentangling the effect of education from the effect of religious homogamy, is to control for the religious composition of the parent couple. However, the correlations mentioned above persist even after adding these controls. According to the estimates from the last model specification (which also includes the child's year of birth as control), for instance, a father with a tertiary education is 4 p.p. less likely to pass on his religion than a father with primary education, while a mother with a tertiary education is 5 p.p. more likely to pass on hers than a mother with primary education.

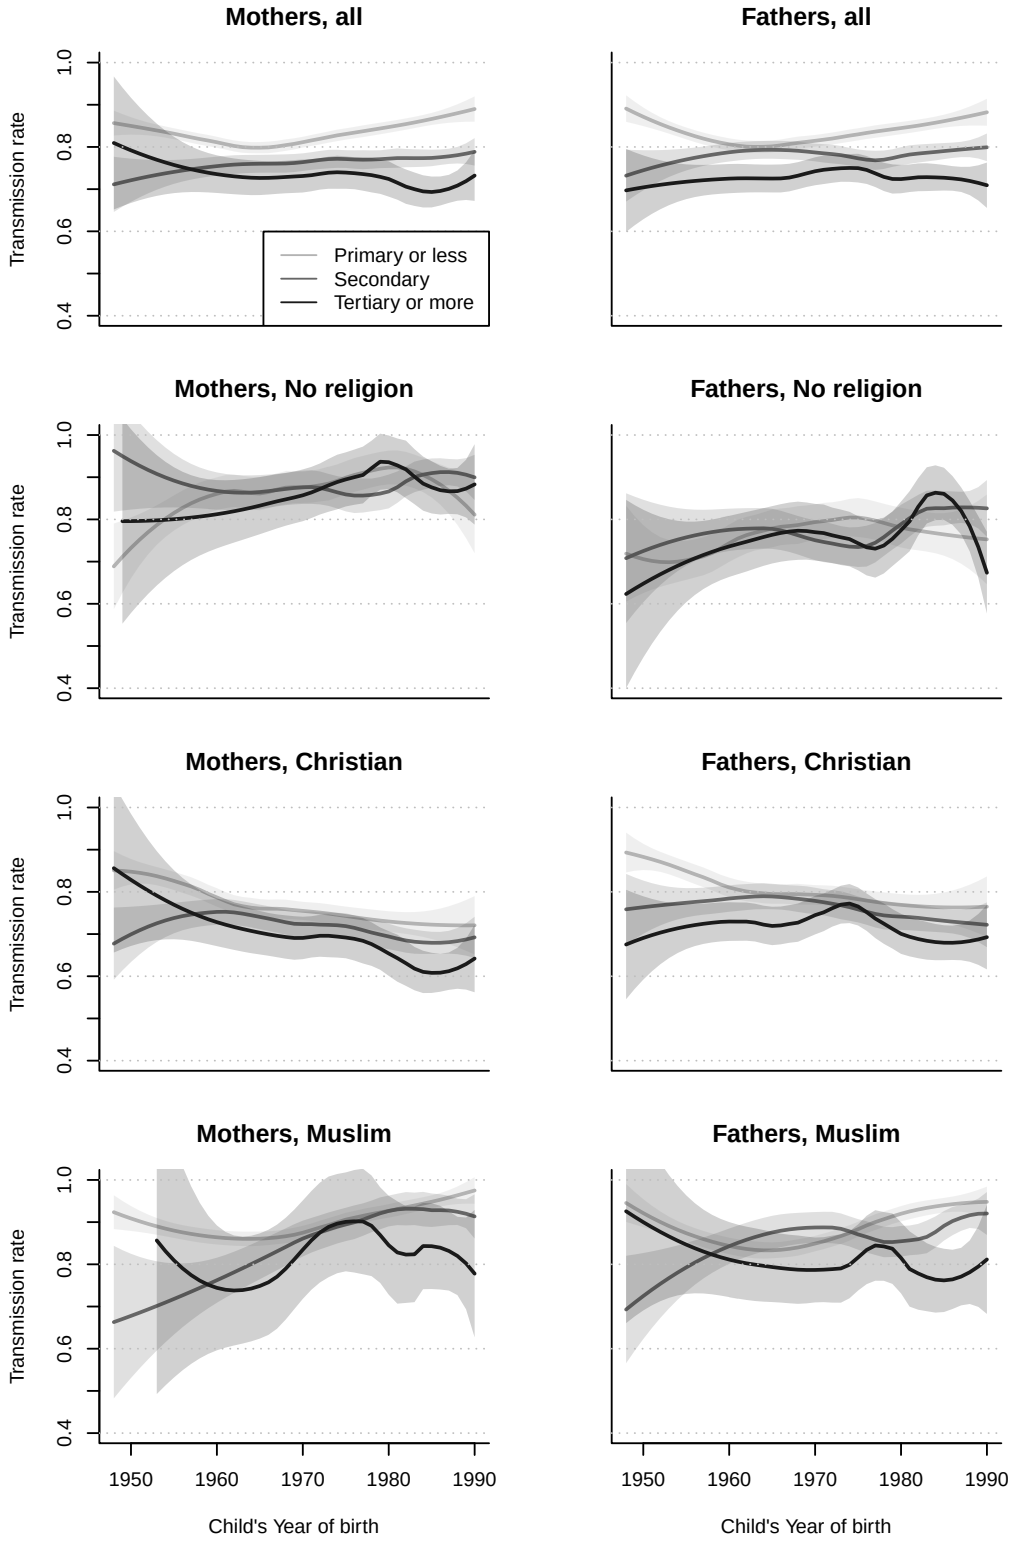

Figure 1.13: Religious transmission by mothers and fathers, by Education.

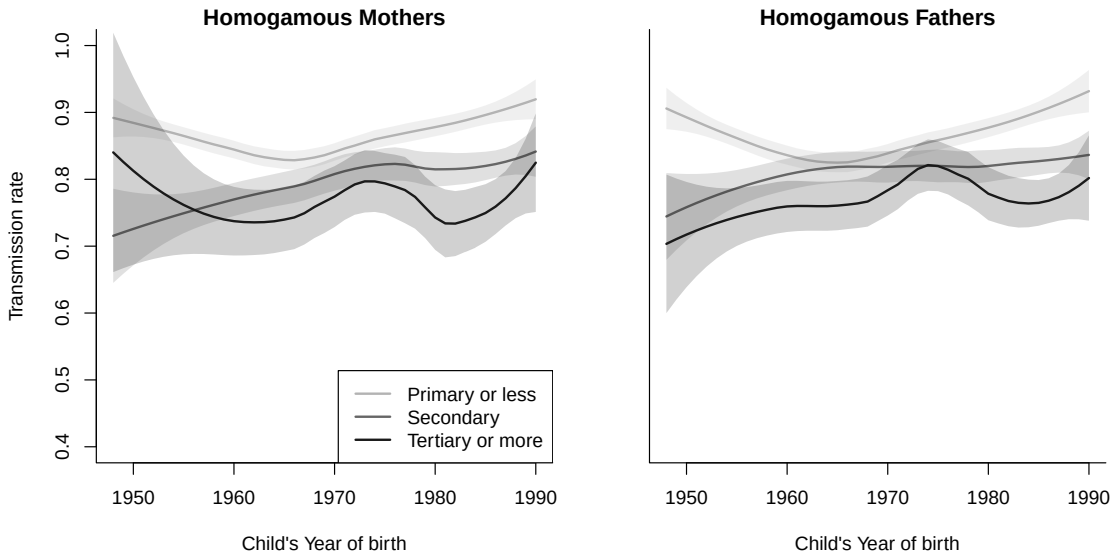

Figure 1.14: Religious transmission by homogamous mothers and fathers, by Education and excluding Nones.

1.3 Reduced-form analysis

The descriptive analysis of section 2.3 yielded several useful insights regarding the intergenerational transmission of religious affiliations. First, vertical transmission from parents to children is very strong, with approximately 87% of respondents in the sample sharing the religious affiliation of at least one parent (Figure 1.11). Second, there is strong heterogeneity in transmission patterns across parental religious affiliations and genders (Figure 1.12). Third, parental education seems to have a detrimental overall effect on religious transmission (Figure 1.13). However, it remains unclear to what extent these differences are driven by marriage patterns rather than by different contributions to socialization across affiliations and genders. For instance, Muslims transmit more on average than Christians or Nones (Figure 1.12), but they also marry more homogamously (Figure 1.5). In this case, is the higher transmission rate of Muslims driven by higher contributions to their children's religious socialization, or by higher homogamy? Can we quantify these contributions?

In this section I turn to a reduced-form model to analyze the transmission of religious affiliation in the TeO data. Religious affiliation is a discrete trait, and I choose a multinomial logit model to investigate its determinants (McFadden 1973). Following the theory of Bisin and Verdier (2000) and subsequent empirical work on cultural transmis-

sion,[9] I focus on two main predictors for the probability that an individual will report a given religious affiliation: the affiliations of her parents, and the shares of each religion in the population. This approach allows me to quantify the importance of parental contributions (vertical socialization) according to their gender, religious affiliation, and educational attainment; versus the role played by the environment (oblique socialization) in determining respondents' affiliations. Furthermore, controlling for parents' religious affiliations in the transmission process also allows to disentangle the effect of homogamy from that of heterogeneous parental contributions.

This section is organized into two parts. In section 1.3.1 I introduce and estimate the baseline econometric specification (a multinomial logit model). The results suggest that vertical socialization plays a more important role than oblique socialization in the transmission process. Furthermore, within vertical socialization, mothers contribute more than fathers, and religious minorities contribute more than religious majorities. Overall, I find that this reduced-form model is very efficient at predicting religious transmission patterns, thereby confirming that parents' religion is a very powerful predictor of the child's religion. In section 1.3.2 I refine the specification to focus on the role of parental education in the transmission process. My results suggest that effects are heterogeneous across religious affiliations, but the general trend indicates that parental contributions to religious socialization rather decrease with their education level.

1.3.1 Multi-logit transmission

Econometric model. Consider a sample of individuals indexed by i, each of whom ultimately chooses one religious affiliation among N available. The propensity for individual i to choose religious affiliation n is measured by a latent variable, which can be written as the product of two components: a component K_{in}, which depends on her observable characteristics, and a component ξ_{in}, which is random. Prefiguring the model of section 1.4.1, I call the observable component K_{in} the *religious capital* of individual i in religion n.

As is standard in discrete choice models, assume now that individual i ultimately

[9] See for instance Bisin and Topa (2003), Bisin, Topa and Verdier (2004), Patacchini and Zenou (2016), Bisin and Tura (2022).

chooses to report the affiliation associated with the largest value among all latent variables:

$$\arg\max_{n} \; K_{in} \times \xi_{in}.$$

If the ξ_{in} are i.i.d. Fréchet (that is, if their logarithms $\ln(\xi_{in})$ are i.i.d. Gumbel), then the probability that i will choose affiliation n is

$$\pi_{in} = \frac{K_{in}}{\sum_{\ell=1}^{N} K_{i\ell}} = \frac{\exp(\ln K_{in})}{\sum_{\ell=1}^{N} \exp(\ln K_{i\ell})}, \tag{1.1}$$

where the second expression makes explicit the link with the multilogit model by using the standard softmax function (generalization of the logistic function to multiple dimensions). Hence, log-religious capital $\ln(K_{in})$ plays a role equivalent to mean utility in the usual discrete choice with random utility framework.

To complete the multilogit model, I must select an econometric specification for the log-religious capital $\ln(K_{in})$ as a function of the observable individual characteristics. I consider a simple model in which the propensity for individual i to choose affiliation n depends on whether her parents have affiliation n and on the share of affiliation n in her environment. This choice can be understood as a broad interpretation of the Bisin and Verdier (2000) cultural transmission model, in which transmission is carried out by parents and by role models outside the family (vertical and oblique socialization). As a starting point, I suppose that parental contributions to religious socialization depend only on their gender and religious affiliation. Thus, mothers n (resp. fathers n) provide a fixed contribution m_n (resp. f_n) toward individuals' propensity to choose affiliation n.

To capture the influence of the environment, I use religions' population shares q_{in} as a proxy. Note that these population shares q_{in} are individual-specific, reflecting that different individuals may be socialized in different cultural environments. Ideally, one could exploit individual variation in two dimensions to explain these differences in oblique socialization. First, the individual's geographical location: the religious mix varies locally, leading to different patterns of oblique socialization. Second, the individual's date of birth: the religious environment has also evolved with time. Unfortunately, religions' population shares in France are comprehensively available neither at the local level nor across time. On locality, the available data is insufficient to obtain credible measures: this would require a dense, large-scale collection of individual religious affiliation in France (which

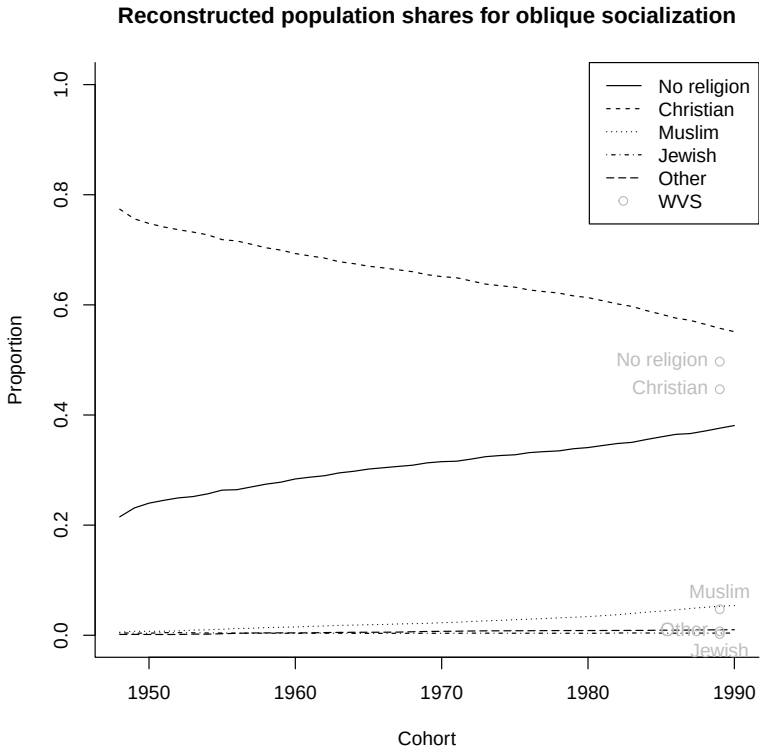

Figure 1.15: Religions' population shares, reconstructed from the TeO survey. Comparison points are taken from the World Values Survey (5th wave, 2005–2009).

is prohibited by law) or, for instance, a comprehensive survey of places of worship of all religions across the country. On time variation however, the TeO data is sufficiently dense to build a credible measure of religious shares in the country across the period of interest. In Appendix 1.7.2, I explain how I reconstruct such a time series of religions' population shares in France: the idea is to consider the religious shares in the subsample of individuals (respondents and their parents) who were alive in a given year. These reconstructed population shares, from 1948 to 1990, are presented in Figure 1.15. In practice, for q_{in} I use the countrywide population shares corresponding to the year in which individual i turned 18 years old.

To summarize, I use the following econometric specification:

$$\ln K_{in} = k_n + m_n \mathbf{1}_{\{i\text{'s mother is } n\}} + f_n \mathbf{1}_{\{i\text{'s father is } n\}} + \alpha\, q_{in}. \tag{1.2}$$

In this expression, the parameters to estimate are k_n, m_n, f_n (for each religion n), and α. I have already mentioned that m_n and f_n correspond to the contributions to religious socialization by mothers n and fathers n respectively. In addition, α measures the impor-

tance of oblique socialization. Finally, the constant k_n captures religion-specific effects in the socialization process. In the abstract, k_n measures the probability of an individual reporting the religious affiliation n in the hypothetical scenario in which she would not have received any socialization, vertical or oblique. In practice, a higher k_n may reflect that the religious affiliation n demands little in terms of knowledge of its affiliates; or that it makes particular efforts to gain new affiliates (beyond the role played by its population share). For this reason, we can expect the "No religion" affiliation to have a high k_n because, by definition, it requires little if any active teaching. In contrast, we can expect the Jewish affiliation to have a low k_n, because it is mostly passed on vertically from the mother.

If we gather all these parameters into a vector $\boldsymbol{\beta}$ and define vectors of individual characteristics $\mathbf{z}_{in}$ appropriately, we can rewrite $\ln(K_{in})$ concisely as

$$\ln K_{in} = \mathbf{z}_{in} \cdot \boldsymbol{\beta}.$$

Thus, equations (1.1) and (1.2) together define a conditional logit model (McFadden 1973, Greene 2008). The conditional logit structure implies that all the parameters m_n and f_n are identified, unlike in the more standard multinomial logit where they are only determined up to a constant. This is because the model (1.2) imposes restrictions compared to a standard multinomial logit model. Specifically, in a multinomial logit model the variable $\mathbf{1}_{\{i\text{'s mother is } n\}}$ would be allowed to have an effect on any latent variable predictor $\ln(K_{i\ell})$; here this effect is assumed to be zero if $\ell \neq n$. The same can be said for the variables $\mathbf{1}_{\{i\text{'s father is } n\}}$ and q_{in}, which have no effect on $\ln(K_{i\ell})$ if $\ell \neq n$. In contrast, the parameters k_n are identified only up to an additive constant.

Testable restrictions. The model imposes restrictions on the transmission probabilities. These restrictions ultimately originate from the independence of irrelevant alternatives assumption inherent to the conditional logit model,

$$\ln\left(\frac{\pi_{in}}{\pi_{i\ell}}\right) = (\mathbf{z}_{in} - \mathbf{z}_{i\ell}) \cdot \boldsymbol{\beta} \qquad (\forall n, \ell). \tag{1.3}$$

Call $\pi_{in\,|\,yab}$ the probability that an individual i will acquire trait n conditional on belonging to the cohort y, and having a mother a and a father b. We can use the last expression

to show (see Appendix 1.7.2) that (1.3) implies

$$\ln\left(\frac{\pi_{ia\,|\,yaa}}{\pi_{ib\,|\,yaa}}\right) - \ln\left(\frac{\pi_{ia\,|\,yab}}{\pi_{ib\,|\,yab}}\right) - \ln\left(\frac{\pi_{ia\,|\,\tilde{y}ba}}{\pi_{ib\,|\,\tilde{y}ba}}\right) + \ln\left(\frac{\pi_{ia\,|\,\tilde{y}bb}}{\pi_{ib\,|\,\tilde{y}bb}}\right) = 0 \quad (\forall a, b, y, \tilde{y}). \tag{1.4}$$

The issue with formally testing this equality however, is that many of these cells (individuals born in year y with a mother a and a father b) have very few or even no observations. For this reason, as an approximation I ignore the role of cohorts y, $\tilde{y}$, and I test whether the equality

$$\ln\left(\frac{\pi_{ia\,|\,aa}}{\pi_{ib\,|\,aa}}\right) - \ln\left(\frac{\pi_{ia\,|\,ab}}{\pi_{ib\,|\,ab}}\right) - \ln\left(\frac{\pi_{ia\,|\,ba}}{\pi_{ib\,|\,ba}}\right) + \ln\left(\frac{\pi_{ia\,|\,bb}}{\pi_{ib\,|\,bb}}\right) = 0 \qquad (\forall a, b) \tag{1.5}$$

holds in the sample, where $\pi_{in\,|\,ab}$ is the probability that an individual i will acquire trait n conditional on having a mother a and a father b (but no longer conditional on the birth cohort y). This simplification relies on the assumption that the population shares q_{in} are not moving drastically over the period considered (see Figure 1.15). In total there are $N^2 = 25$ such tests to perform. Those for which $a = b$ are trivially verified. Those for which $b > a$ have a symmetric equivalent with $a > b$. This leaves 10 tests to perform. Estimators for $\pi_{in\,|\,ab}$ follow binomial distributions, which I use to construct 95% confidence intervals through simulation (parametric bootstrap; see Appendix 1.7.2 for details). Among these 10 tests, 5 tests cannot reject the null hypothesis that (1.5) holds (None–Christian, None–Jewish, None–Other, Christian–Muslim, Muslim–Other); 3 tests reject the null hypothesis (None–Muslim, Christian–Jewish, and Christian–Other); and the 2 other tests cannot be computed due to lack of observations. Overall, Bonferroni's method for global testing rejects (1.5), and the multiple testing procedure by Benjamini and Hochberg (1995) leads to the rejection of the same 3 hypotheses as the separate individual tests (see Appendix 1.7.2).

We cannot rule out that this rejection stems from ignoring cohort effects (i.e. testing (1.5) instead of (1.4)) – in other words, that it arises ultimately from the sparseness of observations along the dimensions considered. Similarly, one would ideally want to test this hypothesis with population shares q_{in} that vary not only across time but also across locality. Nevertheless, rejection of the restriction (1.5) might also warrant the inclusion of other explanatory variables in the econometric specification (1.2). I tackle this issue below by adding an interaction term to the model (1.2), and by considering the effect of

parental education on socialization contributions.

Estimation. We can now proceed with the estimation of the model defined by equations (1.1) and (1.2). Identification comes from the variation in the respondents' religious affiliation. The mothers' contributions to socialization m_n are identified through variation in the father's religion; the fathers' contributions f_n are identified symmetrically; and the oblique socialization coefficient α is identified through cohort variation in population shares. As in a multinomial logit, the intercepts k_n are only identified up to a constant: I choose the most common affiliation, Christian, as the baseline category. With $N = 5$ traits under consideration (None, Christian, Muslim, Jewish, and Other), this leaves a total of $3N - 1 + 1 = 15$ free parameters. I estimate $\boldsymbol{\beta}$ by maximum likelihood, where the log-likelihood is

$$\begin{aligned} \ln L &= \sum_i w_i \sum_{n=1}^{N} \mathbf{1}_{\{i \text{ is } n\}} \times \ln \pi_{in} \\ &= \sum_i w_i \left[\left(\sum_{n=1}^{N} \mathbf{1}_{\{i \text{ is } n\}} \ln K_{in} \right) - \ln \left(\sum_{\ell=1}^{N} K_{i\ell} \right) \right], \end{aligned} \tag{1.6}$$

where the w_i are probabilistic sampling weights provided in the TeO survey. For the covariance matrix I compute the BHHH estimator (Berndt, Hall, Hall and Hausman 1974), from which I obtain the standard errors reported throughout this section.

The results are presented in Table 1.4, column 1. First, consider the parental contributions m_n and f_n, measuring vertical socialization. These parental contributions are highest among minorities (Muslims and Jews, and to a lesser extent, Others), suggesting that the cultural substitution property proposed by Bisin and Verdier (2000) holds here. They are lower for Christians, and close to zero for Nones. Comparing maternal and paternal contributions within a given affiliation, we see that Jewish mothers make significantly higher contributions than Jewish fathers: this result is consistent with the fact that being Jewish is transmitted primarily through the mother. Mothers also contribute more than fathers among Nones and Others, although the difference is less striking. Finally, among Muslims and Christians, mothers and fathers contribute almost equally. Second, the magnitude of oblique socialization is comparable to but less than that of vertical socialization. The estimate for α implies, for instance, that a 50% population share induces an oblique socialization contribution equivalent to half the contribution of

a Christian mother, or one-quarter of the contribution of a Jewish father. Third and last, the estimates for the intercepts k_n can be interpreted in light of the specificities of each affiliation. The intercept for None is the highest, reflecting the fact that while actual religions need to be taught, being nonreligious can simply result from the absence of any religious teaching. This characteristic makes the "No religion" trait special, as it can be acquired not only through active socialization (to secularism, atheism) but also through the lack of socialization in other religions. For this reason, it makes sense that transmission is biased by default toward the "No religion" trait. In the French context specifically, this bias could also account for the socialization influence of schools, which are mostly secular. The intercept for Jewish, in contrast, is significantly lower than the others, so that individuals are very unlikely to become Jewish unless they have a Jewish parent. This result is consistent with the fact that Judaism is not a proselytic religion and is mostly transmitted from parents to children. The intercepts for Muslim and Other are not significantly different from the Christian reference category, which suggests that no stark structural difference exists in the way these religions are transmitted.

Model fit. The specification (1.2) can be compared to the null model defined by an intercept only, $\ln(K_{in}) = k_n$, which has deviance 51 268. With an LR test statistic of $51\,268 - 20\,948 = 30\,320$ on 11 degrees of freedom (which is significant at any conventional confidence level), the model (1.2) explains the data significantly better than the null model. The associated pseudo-R^2 is 0.46, also indicating a good model fit.

Using the estimated parameters, I simulate transmission rates to see how well the model fits aggregate patterns in the data. Figure 1.16 presents observed vs. estimated transmission rates for the three main religious affiliations. Overall, the estimated rates very closely match the observed rates. Sharp turns in the observed transmission rate which are due to cohort variations in parental homogamy rates (e.g. in 1955 for Muslim fathers) are even well replicated by the simulations, thus suggesting that the model indeed manages to disentangle the effect of homogamy from that of parental contributions.

Table 1.4: Conditional logit transmission, estimates.

	Conditional logit estimates	
	model (1.2)	model (1.7)
Constant k_n		
None	2.86^{***} (0.08)	2.76^{***} (0.08)
Christian	0 (baseline)	0 (baseline)
Muslim	-1.32^{***} (0.08)	-1.58^{***} (0.09)
Jewish	-2.93^{***} (0.11)	-3.00^{***} (0.16)
Other	-0.66^{***} (0.08)	-0.79^{***} (0.09)
Mother's contribution m_n		
None	0.11 (0.08)	-0.61^{***} (0.13)
Christian	2.24^{***} (0.08)	1.92^{***} (0.10)
Muslim	3.80^{***} (0.21)	4.67^{***} (0.33)
Jewish	5.33^{***} (0.28)	5.17^{***} (0.32)
Other	3.90^{***} (0.13)	3.85^{***} (0.13)
Father's contribution f_n		
None	0.30^{***} (0.07)	0.05 (0.08)
Christian	1.30^{***} (0.07)	0.62^{***} (0.14)
Muslim	3.22^{***} (0.22)	3.66^{***} (0.32)
Jewish	3.45^{***} (0.43)	2.77^{**} (1.31)
Other	0.78^{***} (0.18)	1.51 (1.06)
Interaction contribution b_n		
None		1.10^{***} (0.16)
Christian		0.90^{***} (0.16)
Muslim		-1.20^{**} (0.47)
Jewish		0.93 (1.39)
Other		-0.68 (1.06)
Oblique socialization coefficient α	1.36^{***} (0.08)	1.37^{***} (0.08)
Observations	20 547	20 547
Sampling weights	Yes	Yes
Deviance $(-2 \ln L)$	21 937	21 901
Pseudo-R^2	0.46	0.46
LR test p-value	baseline	0.000

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

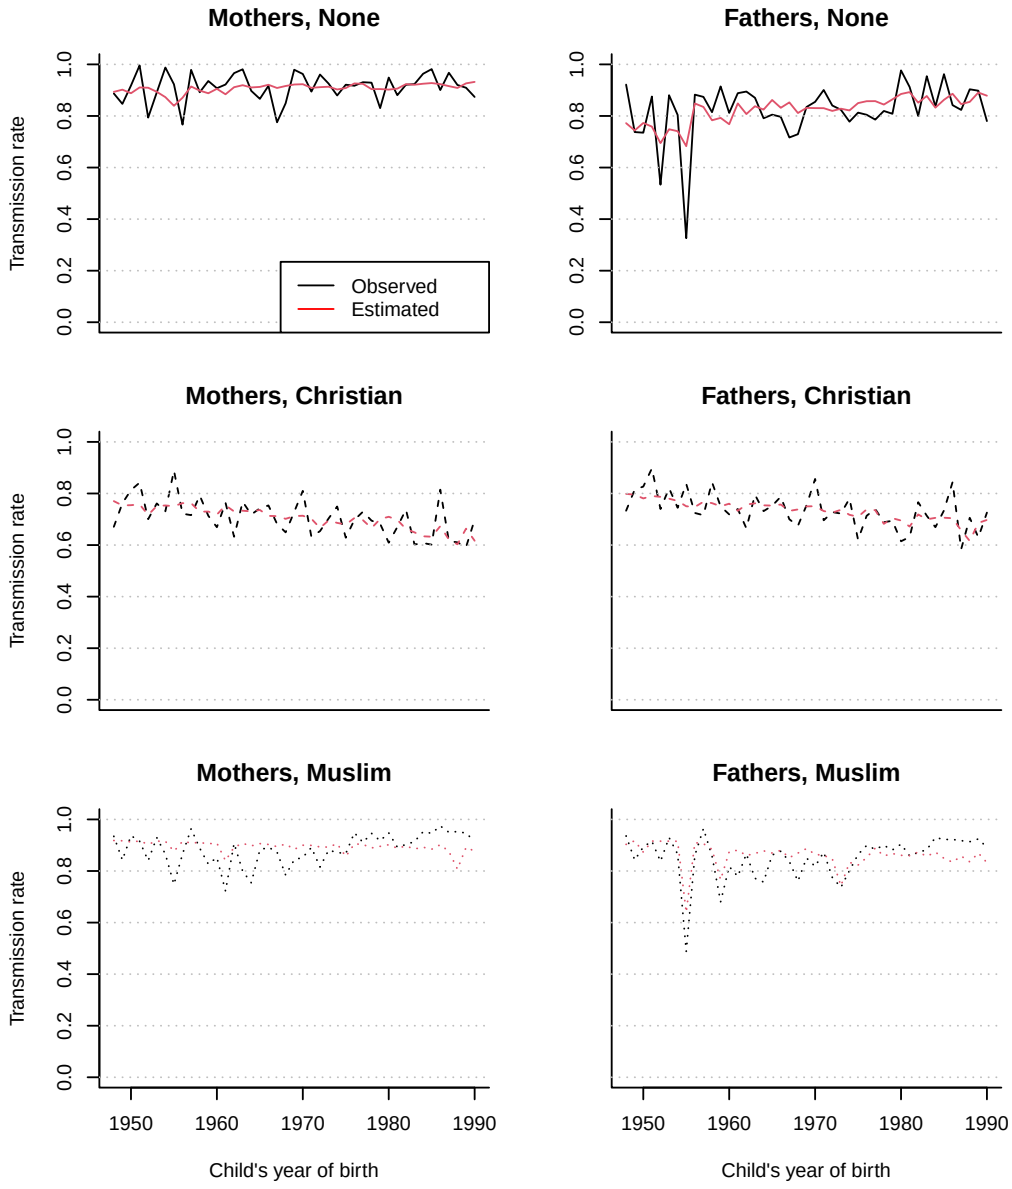

Figure 1.16: Conditional logit transmission, observed vs. estimated transmission rates (by Religion).

Complementarities. The baseline model (1.2) rules out complementarities between the affiliations of the parents. I address this by adding interaction effects to the model,

$$\ln K_{in} = k_n + m_n \mathbf{1}_{\{i\text{'s mother is } n\}} + f_n \mathbf{1}_{\{i\text{'s father is } n\}} + b_n \mathbf{1}_{\{i\text{'s mother is } n\}} \times \mathbf{1}_{\{i\text{'s father is } n\}} + \alpha\, q_{in}, \quad (1.7)$$

so that b_n measures the additional effect of having both parents of religion n on the log-religious capital $\ln(K_{in})$.

The estimation results are presented in Table 1.4, column 2. Compared to the model without interaction effects, the likelihood ratio test statistic is $21\,937 - 20901 = 36$ on 5 degrees of freedom (p-value $\simeq 10^{-6}$), validating the inclusion of these interaction terms as relevant predictors. The interaction parameters are positive for Nones, Christians, and Jews; and negative for Muslims and Others. However, the most precise estimates are for Nones and Christians, pointing toward a complementarity of the parents' religious affiliations in their socialization contributions.

1.3.2 Religious socialization and parental education

To learn more about the potential effect of parents' education levels on the transmission of religion, I extend the previous model by allowing socialization contributions to differ across education levels. Suppose that contributions to socialization now depend not only on the parent's religion n but also on their education level $e \in \{1, \dots, E\}$. Mothers' contributions are thus denoted m_{ne} and fathers' f_{ne}, with ne the bidimensional trait {religion, education} of the parent. The religious capital of i in trait n is now predicted by the following equation:

$$\ln K_{in} = k_n + \sum_e (m_{ne} \mathbf{1}_{\{i\text{'s mother is } ne\}} + f_{ne} \mathbf{1}_{\{i\text{'s father is } ne\}}) + \alpha\, q_{in}. \quad (1.8)$$

The requirement that parents' educational attainments be known in addition to their religious affiliation leads to some sample attrition, down to 18 155 observations from 20 547 previously. As a baseline, I re-estimate the specification (1.2) on this subsample (Table 1.5, column 1). The estimates do not vary significantly from those obtained when using the full sample.

Table 1.5: Conditional logit transmission with education effects, Estimates.

	model (1.2)	model (1.8) $E = \{1, 2, 3\}$	model (1.8) $E = \{1, 2 \text{ or } 3\}$	model (1.8) $E = \{1 \text{ or } 2, 3\}$
Constant k_n				
None	2.82 (0.09)	2.80 (0.14)	2.76 (0.14)	2.83 (0.13)
Christian	0.00 (base)	0.00 (base)	0.00 (base)	0.00 (base)
Muslim	–1.37 (0.09)	–1.42 (0.21)	–1.48 (0.21)	–1.35 (0.20)
Jewish	–2.72 (0.12)	–2.83 (0.47)	–2.85 (0.44)	–2.76 (0.46)
Other	–0.72 (0.09)	–0.76 (0.22)	–0.83 (0.22)	–0.70 (0.21)
Mother's contributions m_{ne}				
None	0.07 (0.09)			
Primary or less		–0.09 (0.13)	–0.11 (0.13)	
Secondary (or more / or less)		0.16 (0.14)	0.22 (0.12)	0.03 (0.11)
Tertiary or more		0.36 (0.19)		0.35 (0.19)
Christian	2.26 (0.09)			
Primary or less		2.21 (0.11)	2.23 (0.11)	
Secondary (or more / or less)		2.23 (0.11)	2.27 (0.11)	2.22 (0.10)
Tertiary or more		2.52 (0.12)		2.52 (0.12)
Muslim	3.82 (0.23)			
Primary or less		3.80 (0.13)	3.79 (0.13)	
Secondary (or more / or less)		4.26 (0.16)	3.98 (0.15)	3.90 (0.12)
Tertiary or more		2.74 (0.28)		2.83 (0.28)
Jewish	5.79 (0.32)			
Primary or less		4.71 (0.78)	4.66 (0.79)	
Secondary (or more / or less)		5.58 (0.64)	6.06 (0.49)	5.26 (0.55)
Tertiary or more		6.53 (0.80)		6.52 (0.79)
Other	3.76 (0.16)			
Primary or less		3.97 (0.18)	3.97 (0.18)	
Secondary (or more / or less)		3.21 (0.20)	3.23 (0.19)	3.78 (0.17)
Tertiary or more		3.27 (0.28)		3.45 (0.28)
Father's contributions f_{ne}				
None	0.39 (0.07)			
Primary or less		0.41 (0.11)	0.41 (0.11)	
Secondary (or more / or less)		0.34 (0.11)	0.38 (0.10)	0.37 (0.10)
Tertiary or more		0.59 (0.14)		0.60 (0.14)
Christian	1.24 (0.07)			
Primary or less		1.32 (0.10)	1.31 (0.10)	
Secondary (or more / or less)		1.23 (0.10)	1.18 (0.09)	1.28 (0.09)
Tertiary or more		1.00 (0.11)		1.01 (0.11)
Muslim	3.28 (0.23)			
Primary or less		3.19 (0.14)	3.20 (0.14)	
Secondary (or more / or less)		3.54 (0.16)	3.48 (0.15)	3.22 (0.14)
Tertiary or more		3.70 (0.21)		3.69 (0.21)
Jewish	3.04 (0.48)			
Primary or less		3.34 (0.78)	3.24 (0.80)	
Secondary (or more / or less)		2.86 (0.72)	3.39 (0.51)	3.06 (0.52)
Tertiary or more		5.65 (0.79)		5.65 (0.79)
Other	1.13 (0.23)			
Primary or less		0.86 (0.19)	0.86 (0.19)	
Secondary (or more / or less)		1.20 (0.19)	1.53 (0.17)	1.09 (0.17)
Tertiary or more		1.60 (0.21)		1.34 (0.21)
Oblique socialization coefficient α	1.41 (0.08)	1.37 (0.28)	1.25 (0.28)	1.46 (0.27)
Observations	18 155	18 155	18 155	18 155
Sampling weights	Yes	Yes	Yes	Yes
Deviance $(-2 \ln L)$	19 198	19 139	19 174	19 158
Pseudo-R^2	0.47	0.47	0.47	0.47
LR test p-value	baseline	0.000	0.007	0.000

Note: Standard errors in parentheses.

Primary vs. Secondary vs. Tertiary. The estimation procedure for the model with education effects remains the same, except that we now have $N + 2NE = 35$ free parameters to estimate. First, I consider the three education levels that I used in section 2.3: Primary or less, Secondary, and Tertiary or more. The results are presented in Table 1.5, column 2. The likelihood ratio test statistic is $19\,198 - 19\,139 = 59$ on 20 degrees of freedom for a p-value smaller than 10^{-5}, providing evidence that the parents' education levels matter in predicting transmission rates. Regarding the estimates, there is no clear pattern for the relationship between education and socialization contributions. The estimated parameters remain qualitatively close to those estimated in the model without educational effects.

Primary vs. Secondary. I attempt to estimate the effect of education more precisely by reducing the number of educational categories to two: Primary or less ($e = 1$), and Secondary or more ($e = 2$). The results are presented in Table 1.5, column 3. The likelihood ratio test statistic is $19\,198 - 19\,174 = 24$ on 10 degrees of freedom, for a p-value of 0.007. Once again, estimated contributions are qualitatively close to their level in the absence of education effects.

With only two education levels, it is also easier to verify whether human capital has a discernable effect on socialization contributions. To do so, I test the two hypotheses $m_{n1} = m_{n2}$ and $f_{n1} = f_{n2}$ for every religion n. I recover the distributions of $\hat{m}_{n1} - \hat{m}_{n2}$ and $\hat{f}_{n1} - \hat{f}_{n2}$ by the delta method, which I use to construct 95% confidence intervals for $m_{n1} - m_{n2}$ and $f_{n1} - f_{n2}$ (Figure 1.17, left panel). Differences between the socialization contributions of Primary or less parents and Secondary or more parents are not statistically significant at the 5% level, except for Other parents. This result is consistent with the estimates from the linear model of Table 1.3 in the descriptives, which showed that a Secondary education had little to no significant effect on transmission rates for either mothers or fathers. Furthermore, note that among Others education has opposite effects on socialization contributions for mothers and fathers.

Secondary vs. Tertiary. Finally, I replicate this exercise with the following two education levels: Secondary or less, and Tertiary or more (Table 1.5, column 4). The likelihood ratio test statistic is $19\,198 - 19\,158 = 40$ on 10 degrees of freedom, for a p-value close to 10^{-5}. Again, I test whether education has a significant effect on socialization contribu-

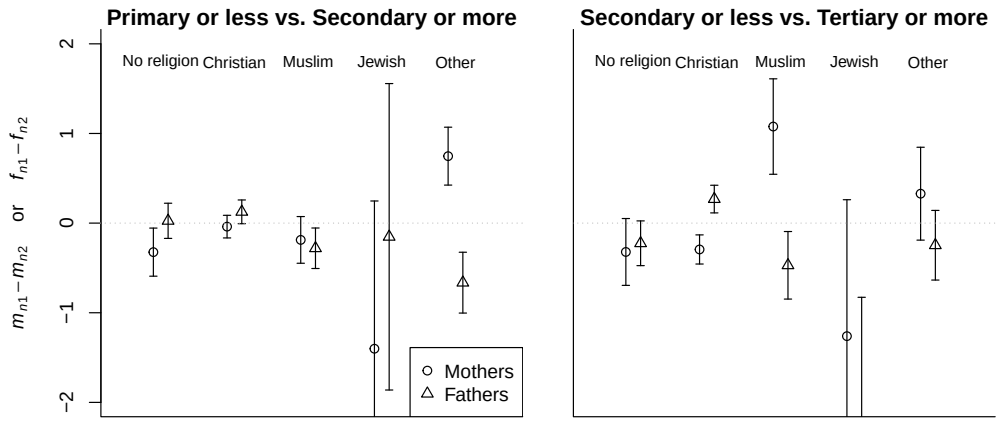

Figure 1.17: Differences in socialization contributions across education levels. 95% confidence intervals for $m_{n1} - m_{n2}$ and $f_{n1} - f_{n2}$ are reported for each religion n. Left panel: Primary or less ($e = 1$) vs. Secondary or more ($e = 2$). Right panel: Secondary or less ($e = 1$) vs. Tertiary or more ($e = 2$).

tions for mothers and fathers of all religions (Figure 1.17, right panel). However, the effect of education remains heterogeneous across mothers and fathers, and across religions.

It is difficult to say whether these differences across religions are structural or if they are the result of a model misspecification. The education of the parents could impact their opportunity cost of socializing their child. It could also shift the power balance in the couple, or be associated with different preferences for religious socialization. The estimation of the structural model in section 1.5 will allow me to shed some light on the possible mechanisms at play.

1.3.3 Alternative model for the influence of education

Suppose instead that education has a uniform multiplicative effect on contributions across religious affiliations, so that

$$\begin{aligned}\ln K_{in} = k_n &+ (1 + \kappa_1 \mathbf{1}_{\{i\text{'s mother has } e \geq 2\}} + \kappa_2 \mathbf{1}_{\{i\text{'s mother has } e \geq 3\}}) \times m_n \mathbf{1}_{\{i\text{'s mother is } n\}} \\ &+ (1 + \rho_1 \mathbf{1}_{\{i\text{'s father has } e \geq 2\}} + \rho_2 \mathbf{1}_{\{i\text{'s father has } e \geq 3\}}) \times f_n \mathbf{1}_{\{i\text{'s father is } n\}} + \alpha\, q_{in}. \quad (1.9)\end{aligned}$$

This model is still linear in the observables, but it imposes more structure than the previous model from section 1.3.2. Indeed, here I impose $m_{n2} = (1 + \kappa_1)m_{n1}$ and $m_{n3} = (1 + \kappa_1 + \kappa_2)m_{n1}$ for all religions n, where the parameters κ_1 and κ_2 do not depend on n (and similarly for the f_{ne} with ρ_1 and ρ_2). The goal is to observe the effects of education

by measuring κ_1, κ_2, ρ_1, and ρ_2. Negative values, for instance, would provide evidence of lower average contributions for higher-educated parents.

The results are presented in Table 1.6. In the first column I re-estimate the baseline model on the subsample of individuals for whom the educational attainment of both parents is available. The estimates remain comparable to those from the full sample estimation. In the second column, I estimate the new model with multiplicative education effects. The estimates for κ_1, κ_2, ρ_1, and ρ_2 suggest that education has opposite effects on contributions to socialization across genders: positive for mothers and negative for fathers. For mothers, more education is associated with higher contributions: mothers with a Tertiary education or more make contributions that are 10% higher than mothers who have a Primary education or less. For fathers it is the opposite: a Tertiary education is associated with contributions that are 13% lower. The effect of having a Secondary education goes in the same direction (1% higher contributions for mothers, 2% lower for fathers) but is not statistically significant. The LR test value is 10 on 4 degrees of freedom (p-value = 0.033), so the added parameters provide significant explanatory power to the model.

Table 1.6: Conditional logit transmission with education effects, estimates.

	Conditional logit estimates	
	(1.2)	(1.9)
Constant k_n		
None	2.82*** (0.09)	2.87*** (0.09)
Christian	0 (baseline)	0 (baseline)
Muslim	−1.37*** (0.09)	−1.32*** (0.10)
Jewish	−2.72*** (0.12)	−2.70*** (0.12)
Other	−0.72*** (0.09)	−0.67*** (0.10)
Mother's contribution m_n		
None	0.07 (0.09)	0.14 (0.09)
Christian	2.26*** (0.09)	2.15*** (0.09)
Muslim	3.82*** (0.23)	3.75*** (0.23)
Jewish	5.79*** (0.32)	5.60*** (0.31)
Other	3.76*** (0.16)	3.70*** (0.16)
Father's contribution f_n		
None	0.39*** (0.07)	0.30*** (0.07)
Christian	1.24*** (0.07)	1.37*** (0.07)
Muslim	3.28*** (0.23)	3.38*** (0.24)
Jewish	3.04*** (0.48)	3.17*** (0.49)
Other	1.13*** (0.23)	1.20*** (0.23)
Multiplicative education effets		
Secondary or more mother κ_1		0.01 (0.01)
Tertiary or more mother κ_2		0.09*** (0.02)
Secondary or more father ρ_1		−0.02 (0.02)
Tertiary or more father ρ_2		−0.11*** (0.02)
Oblique socialization coefficient α	1.41*** (0.08)	1.48*** (0.09)
Observations	18 115	18 115
Sampling weights	Yes	Yes
Deviance $(-2 \ln L)$	19 198	19 188
Pseudo-R^2	0.47	0.47
LR test p-value	baseline	0.033

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.4 Structural model

Section 1.3 mostly confirmed the patterns documented by the descriptive analysis of section 2.3: mothers contribute more to the religious socialization of their children than fathers, religious minorities contribute more than majorities, and parents' education levels are relevant predictors of their contributions to religious socialization. This section also provided a first quantified measure of these various heterogeneities in the transmission process. Because, in the data, mothers are less educated than fathers and religious minorities less than majorities, it is tempting to use an economic argument involving education to explain these differences in socialization patterns. The estimates from the reduced-form analysis are indeed broadly consistent with an economic explanation, namely, that higher-educated individuals have a higher opportunity cost of socializing their children and, therefore, they socialize them less. In particular, in light of the possible substitution between culture and education discussed in the introduction, we can wonder whether higher-educated parents reallocate resources from cultural socialization toward formal education because they have a comparative advantage in the latter.

To investigate this potential mechanism, in this section I construct a model of intergenerational cultural socialization and human capital formation. First, I focus on modeling the technology available to parents for the cultural socialization of their child. Similar to Iannaccone (1990) in the case of religion, I take a human capital approach to culture, introducing the notion of *cultural capital* as an intensive and multidimensional measure of culture for individuals.[10] This approach considers that culture is not simply a static affiliation but a gradually built ensemble of knowledge and practices in which individuals can invest. Here, I specifically consider the role of parental time investments in building children's cultural capital: this is the socialization process.

Second, I embed this model of socialization within a collective household framework in which parents care about passing on both human capital and cultural capital to their child. To do so, they can allocate their time between two activities: human capital production, and cultural socialization. The goal of the model is to describe a simple trade-off between these two activities. Crucially, the human capital of the parents is assumed to be productive in the human capital formation of the child but not in cultural socialization.

[10] This terminology has of course a long tradition in sociology (Bourdieu 1979), which was itself influenced by the work of Gary Becker.

Thus by construction, parents with higher human capital have a comparative advantage in human capital formation relative to cultural socialization. Given this advantage, an increase in the parents' human capital will lead to a reallocation of time in favor of the child's human capital formation and at the expense of her cultural socialization. This mechanism also interacts with one of the main ideas developed by Bisin and Verdier (2000) on cultural transmission: cultural minorities must make more effort to transmit their culture than majorities because majorities can rely on the public provision of cultural socialization, or *oblique socialization*. In my framework, such effort happens at the expense of human capital formation, thus creating an imbalance between minorities and majorities: all else equal, minorities devote less time to their child's human capital formation than majorities.

This section is divided into three parts. In section 1.4.1, I present a time allocation theory of cultural socialization. In section 1.4.2, I describe the household's decision framework. Finally, in section 1.4.3, I solve the model and provide a short analysis of the trade-offs involved.

1.4.1 A time allocation theory of socialization

To model socialization within the household, I start with the technology of cultural socialization available to the parents. My theoretical approach is grounded in the seminal work on cultural transmission by Bisin and Verdier (2000). In particular, I adopt the distinction between *vertical socialization*, performed by the parents, and *oblique socialization*, performed by the rest of the population. However, my approach also builds upon this work, most notably by considering a continuous, multidimensional *cultural capital* for the child, and by incorporating insights from the literature on human capital formation.

In their model Bisin and Verdier consider culture as a discrete, exclusive trait that is transmitted probabilistically. Instead, I model the culture transmitted to children as multidimensional and with an intensive measure, and I label this the cultural capital of children. Behind this label is the idea that culture is an example of task-specific human capital (Gibbons and Waldman 2004), an approach that has already been adopted in the economics of religion by Iannaccone (1990). As such, cultural capital associated with Christianity for instance, serves a different purpose from cultural capital associated

with Islam.[11] Furthermore, rather than considering all the different possible channels of cultural capital formation, I focus here on the role of parental time investments in their child's cultural socialization. Keeping religion as an illustrative example, the child's cultural capital is then a measure of the intensity of her socialization to Christianity, Islam, Atheism... This modeling choice proves important in disentangling the different influences involved in the cultural socialization process. It is also particularly convenient for transposition to the empirical analysis – indeed, I will show how it maps naturally to the reduced-form analysis from section 1.3. Finally, because I consider cultural capital as a specific modality of human capital, I adapt existing insights from the literature on children's human capital formation (Del Boca et al. 2014, 2016, Chiappori et al. 2017) and the theory of time allocation (Becker 1965) to represent the production of children's cultural capital. Doing so provides tractable solutions to the collective household model in section 1.4.2 while maintaining the intuitive results that would derive from a more agnostic approach.

The formation of cultural capital. Consider a household formed by two parents, indexed by $i \in \{1, 2\}$. For simplicity, I assume that each household has one child. Parent i possesses a single cultural trait n_i among N possible traits. In this model, parents domestically produce the child's cultural capital by spending time on cultural socialization.

To model the accumulation of the child's cultural capital, I rely on existing results from the literature on human capital formation. Specifically, Cunha and Heckman (2007) and Aizer and Cunha (2012) provide evidence that investments in the human capital of children are dynamic complements, in the sense that existing human capital increases the returns of current investments. (Thus, past investments indirectly increase the returns of current investments, hence the "dynamic" complementarity.) This model feature was for instance adopted in an empirical structural framework by Del Boca et al. (2014, 2016). As a specific type of human capital, it is reasonable to assume that cultural capital is produced similarly. Interpreted here in a continuous time setting, this dynamic complementarity means that the cultural capital returns $\mathrm{d}K$ on a marginal time investment $\mathrm{d}s$

[11] This example also emphasizes a possible complementarity between the different dimensions of cultural capital: as Abrahamic religions, there is significant overlap between Christianity and Islam in terms of religious knowledge or practice. If we consider the example of language, the same could be said about languages that share a common script, vocabulary, or grammatical structure. While I will not consider it here for the sake of simplicity, this complementarity between different dimensions of cultural capital could easily be added to the model, at the cost of additional complexity in the number of parameters.

are proportional to the stock of cultural capital already produced. Following this logic, the law of accumulation of the cultural capital K is

$$\mathrm{d}K = K \times a\,\mathrm{d}s \tag{1.10}$$

where a is a positive parameter denoting the time productivity of the individual.[12] Integrating equation (1.10), we find that the log-cultural capital is produced from the time investment s by a linear technology:

$$\ln K = k + a\,s. \tag{1.11}$$

Equation (1.11) thus describes the accumulation of cultural capital when one individual is involved in the child's socialization. In reality however, the cultural socialization of children involves several individuals, most notably the parents. Here, I follow the literature on cultural transmission by assuming that the child is subject to both *vertical* and *oblique* socialization (Bisin and Verdier 2000, 2011). Vertical socialization, on the one hand, results from "purposeful socialization decisions inside the family." In my time allocation framework, it is carried out in the form of (endogenous) parental time inputs s_i spent socializing the child. Oblique socialization, on the other hand, summarizes other socialization processes that happen outside of the family. To model oblique socialization, I first assume that the child has a fixed time attention span for being socialized, which I normalize to 1. Deducting the time taken by the parents leaves time $1 - s_1 - s_2$ during which the child is subject to oblique socialization. Second, I assume that this remaining socialization time is spent randomly with the rest of the population. This means that if a culture has a population share q, the child spends time $(1 - s_1 - s_2)q$ being socialized to that culture. Thus, as in the standard Bisin and Verdier model, oblique socialization to a given culture is proportional to that culture's share in the population: more widespread cultures exert a stronger influence. Accounting for these different socialization channels

[12] This is equivalent to the way that an investment I grows with time under an interest rate r,

$$\mathrm{d}I = I \times r\,\mathrm{d}t \implies I = I_0\,e^{rt}.$$

into account, the child's cultural capital is produced via the technology

$$\ln K = k + a_1\, s_1 + a_2\, s_2 + a_0\,(1 - s_1 - s_2)q \tag{1.12}$$

where a_1 and a_2 are the productivities of parents 1 and 2 respectively, and a_0 is the productivity of oblique socialization.

The socialization technology (1.12) still describes a unidimensional accumulation process. Culture, however, is multidimensional: the child receives socialization in all N cultural traits present in the population. This process constitutes a N-dimensional vector $(K_n)_{1\leq n\leq N}$, where each K_n corresponds to the child's cultural capital in a different trait. The component K_n is increasing in the parental time investments in the child's socialization to trait n, and in the population share q_n of trait n. In the most general case, parents would be able to contribute to the child's socialization to any trait. However, to simplify the analysis, it is useful to consider that a parent can only socialize the child to their own trait. There are at least two reasons to justify this assumption. First, a parent is likely to prioritize transmitting their own culture, and therefore to use their available time doing so. Second, they simply might not have the capacity to transmit another culture if they are not affiliated or familiar with it themselves (e.g. ethnicity, but also language, religion). For this reason, I assume that the time s_i devoted by parent i is fully counted toward the socialization of the child to that parent's trait, n_i. Thus, the child's cultural capital in trait n is formed according to

$$\ln K_n = k_n + a_1\, s_1\, \mathbf{1}_{\{n_1=n\}} + a_2\, s_2\, \mathbf{1}_{\{n_2=n\}} + a_0\,(1 - s_1 - s_2)\, q_n \tag{1.13}$$

where $\mathbf{1}_{\{n_i=n\}}$ is an indicator equal to 1 if and only if parent i has trait n.

Examples. *For fixed parental time inputs s_1, s_2, a child with homogamous parents of culture n will receive the cultural capital*

$$\ln K_n = k_n + a_1\, s_1 + a_2\, s_2 + a_0\,(1 - s_1 - s_2)\, q_n, \qquad \ln K_\ell = k_\ell + a_0\,(1 - s_1 - s_2)\, q_l \quad (\forall \ell \neq n),$$

while a child with heterogamous parents of cultures $n_1 \neq n_2$ will receive

$$\ln K_{n_1} = k_{n_1} + a_1\, s_1 + a_0\,(1 - s_1 - s_2)\, q_{n_1}, \qquad \ln K_{n_2} = k_{n_2} + a_2\, s_2 + a_0\,(1 - s_1 - s_2)\, q_{n_2},$$

$$\ln K_\ell = k_\ell + a_0 \left(1 - s_1 - s_2\right) q_l \quad (\forall \ell \neq n_1, n_2).$$

The model is also readily extendable to single-parent families: for instance, a child with only parent 1 will receive

$$\ln K_{n_1} = k_{n_1} + a_1 \, s_1 + a_0 \left(1 - s_1\right) q_{n_1}, \qquad \ln K_\ell = k_\ell + a_0 \left(1 - s_1\right) q_l \quad (\forall \ell \neq n_1).$$

Below I will introduce the decision framework in which parents choose their time inputs s_i endogenously. We can already imagine, however, how the functional form (1.13) will impact the socialization decisions of the household. First, if the parental productivities a_1 and a_2 are different, one parent has a comparative advantage over the other in the child's socialization. This feature of the model opens up the possibility of productivity-driven specialization in the household, which is one possible way to explain disparities in transmission rates between mothers and fathers. Second, the model assumes that vertical socialization comes at the expense of oblique socialization. Consequently, parents who belong to a more widespread culture have lower returns on the time they spend socializing their children. This point relates to the *cultural substitution* property introduced by Bisin and Verdier (2001), to which I will return during the analysis.

Link with the reduced-form. In the theory above, individuals have a complex cultural identity that is represented by a multidimensional cultural capital. Empirically however, this multidimensional approach to culture can prove problematic. Indeed, to implement this theoretical framework directly with data, the researcher should ideally have an intensive measure of culture along multiple dimensions (e.g. the level of proficiency in several languages). However, in most cases, surveys do not report this kind of measure of the respondents' culture(s). Rather, survey respondents are often categorized into a single, exclusive affiliation (e.g. religion, ethnicity). This is notably the case for the respondents' religious affiliation in the TeO data.

In section 1.3 we have seen that the multinomial logit model addresses this issue by mapping the multidimensional, intensive measure of culture from the theory into an extensive, discrete cultural affiliation as reported in the data. This approach amounts to considering the reporting of a single cultural affiliation as a choice among coexisting cultural identities. Following the discrete choice theory logic, individuals are then more

likely to report a cultural affiliation in which they have higher cultural capital. Moreover, we have seen how the linear form of log-cultural capital naturally fits into a multilogit regression framework.

Note that equation (1.13) bears a striking resemblance to the econometric specification (1.2) used in the reduced-form analysis of section 1.3. This equation provides a theoretical foundation for the log-religious capital $\ln(K_{in})$ being a linear function of the observables $\mathbf{1}_{\{i\text{'s mother is } n\}}$, $\mathbf{1}_{\{i\text{'s father is } n\}}$, and q_{in}. Furthermore, it suggests the use of measures of parental time spent on religious socialization to predict individuals' choice of religious affiliation. Unfortunately, there is no such measure in the TeO data. In its absence, we can interpret the estimated socialization contributions from section 1.3 as proxies of parental socialization time investments. Recall that according to the reduced-form estimates, religious minorities contribute more to the socialization of their children than majorities, and mothers contribute more than fathers. The model can rationalize the difference between mothers and fathers in two ways: mothers are more productive at socialization ($a_1 > a_2$), or they simply spend more time on religious socialization than fathers ($s_1 > s_2$).[13] In contrast, the model can only rationalize the difference between religious minorities and majorities through higher socialization time investments on the part of minorities. Because the econometric specification (1.2) ignored the adverse role of vertical socialization on oblique socialization present in the model (1.13), there is no such direct interpretation of the oblique socialization coefficient. Finally, the constant k_n from the reduced-form analysis could be understood as a measure of the initial stock of religious capital across the different religious affiliations.

Decreasing returns to socialization. For simplicity of exposition, I have assumed that socializing individuals have a constant productivity of socialization, equal to a_1, a_2, or a_0. In fact, it may be more accurate to assume that the socialization time investments of the parents exhibit decreasing returns, in the sense that the marginal productivity of their time declines as they spend more time socializing the child. (See for instance Chiappori et al. 2017 for children's human capital formation). Such declines could occur because parents eventually run out of new knowledge to transmit, or because children

[13] If we denote by s_{1n} and s_{2n} the socialization time investments of mothers n and fathers n respectively, then m_n and f_n are defined as

$$m_n = a_1\, s_{1n}, \qquad f_n = a_2\, s_{2n}.$$

progressively lose attention when taught by a single teacher.

To account for this possibility, I assume that individuals' socialization productivity decreases with the time s spent socializing the child. To keep the model tractable, I consider that productivity decreases linearly: after having spent time s on socialization, an individual has marginal productivity $a \times (1 - \gamma s)$. Under this assumption, a is the initial socialization productivity at $s = 0$, and γ is a positive parameter representing how quickly productivity declines. (Note also that above $s > 1/\gamma$, socialization becomes counterproductive.) The law of accumulation of cultural capital (1.10) is modified to

$$\mathrm{d}K = K \times a \times (1 - \gamma s)\,\mathrm{d}s. \tag{1.14}$$

By integrating this equation we obtain the total cultural capital output produced from a socialization time investment s,

$$\ln K = a\left(s - \frac{\gamma}{2}s^2\right).$$

Since they spend a significant amount of time with their children, it is reasonable to assume that parents are subject to this decline in socialization productivity. Oblique socialization, in contrast, is by assumption carried out by many different individuals who each spend a marginal amount of time socializing the child. For this reason, the time $1 - s_1 - s_2$ dedicated to oblique socialization still produces cultural capital at a constant rate and does not suffer from a decrease in productivity. To summarize, incorporating decreasing returns in socialization yields the following production function for cultural capital:

$$\ln K_n = k_n + a_1\left(s_1 - \frac{\gamma_1}{2}s_1^2\right)\mathbf{1}_{\{n_1=n\}} + a_2\left(s_2 - \frac{\gamma_2}{2}s_2^2\right)\mathbf{1}_{\{n_2=n\}} + a_0\left(1 - s_1 - s_2\right)q_n. \tag{1.15}$$

Note that compared to equation (1.13), here I also included the constant k_n. This is the functional form that I will use in the household model below and in section 1.5 for the structural econometric model.

1.4.2 Household model

After describing the technology of socialization, I now turn to the trade-offs faced by the parents when choosing their socialization time investments. As mentioned above, parents' human capital will play a role here. In addition to the cultural trait n_i, parent i is now also characterized by a human capital level h_i (continuous). Parents have a fixed time budget, which they must allocate between the production of the child's human capital and cultural capital.

The child's cultural capital is produced from the parents' socialization time inputs s_i according to the technology (1.15). I assume that the child's human capital is produced with a fundamentally similar technology from time inputs t_i of the parents. Unlike for cultural socialization however, I assume that the parental human capital h_i increases the productivity of parent i during human capital production.[14] These two assumptions are consistent with existing models of children's human capital formation (Del Boca et al. 2016, Chiappori et al. 2017). Thus the child's human capital H is produced from parental inputs and characteristics according to

$$\ln H = (b_1 + h_1)\left(t_1 - \frac{\gamma_1}{2}t_1^2\right) + (b_2 + h_2)\left(t_2 - \frac{\gamma_2}{2}t_2^2\right). \tag{1.16}$$

As for a_1 and a_2 in the case of cultural capital production, the parameters b_1 and b_2 denote the baseline productivities of parents 1 and 2 respectively. Note also that I have taken a constant equal to 0 in the production function – this is without loss of generality for the model. (For this reason, one could also add a source of "oblique" production of human capital, without consequence for the model's insights.)

Parents care about their child's human capital and cultural capital. To simplify, I assume that parent i values only the cultural capital of the child in their own trait, n_i. Based on this assumption, I consider a Cobb–Douglas utility for parent i of the following form:

$$u_i = \nu_i \ln(K_{n_i}) + \ln H.$$

The parameter ν_i is an important primitive of the model, representing the value of the child's cultural capital (relative to her human capital) for parent i.

[14] The theoretical results of this section would still hold if, instead of parental human capital having no effect of the production of cultural capital, it simply had a smaller effect.

I use a collective household model (Chiappori 1992) to represent the parents' decision-making, so that parental decisions lead to an outcome on the Pareto frontier of the household. In other words, the intrahousehold decisions must maximize a weighted sum of the parents' utilities:

$$\max_{t_i,s_i} \left\{\mu\, u_1 + u_2 = \mu\, \nu_1 \ln K_{n_1} + \nu_2 \ln K_{n_2} + (\mu+1) \ln H\right\}, \tag{1.17}$$

where μ is the relative power (Pareto weight) of parent 1, fixed exogenously. The constraints concern the time available to the parents: I assume a fixed time budget T_i for parent i, so that the household constraints are

$$t_i + s_i \leq T_i, \qquad i = 1, 2. \tag{1.18}$$

These constraints must be saturated at the optimal time allocation as long as γ is small enough compared to T_i.

Discussion. This framework shares similar features with existing models of cultural transmission and human capital formation. The seminal model of cultural transmission was proposed by Bisin and Verdier (2000), in which parents also care about passing on their culture to their children, and can contribute to their child's cultural socialization to their own traits. The fact that parents might want to transmit a different culture than their own is therefore not considered in their model or in mine. A reason for such a preference could be discrimination against or in favor of a given culture. Such a phenomenon has been documented for instance by Saleh (2018), who shows how differential taxation in medieval Egypt incentivized Coptic Christians (who faced higher taxes) to adopt the Muslim affiliation. Botticini and Eckstein (2007, 2012) also show how economic incentives had an impact on conversions from Judaism to Christianity across history. The crucial feature of the Bisin and Verdier model, namely the substitution between vertical and oblique socialization, is also embedded in my model through the cultural capital production technology (section 1.4.1).

My model also departs from Bisin and Verdier in several ways. First, in their model, parents' efforts to socialize the child have an abstract convex cost. In my model, these efforts are specified as time allocations, which are made at the expense of the child's

human capital production.

Second, in my model parents care about the cultural capital in their own trait, as opposed to the transmission probability of every trait in Bisin and Verdier (2000). The two formulations are in fact theoretically equivalent in cases with two traits (which is the case considered by Bisin and Verdier), but this is no longer true when there are three or more traits.[15] However, fewer theoretical results exist beyond the two-traits case (see Montgomery 2010). Assuming that parents care about cultural capital, not transmission probabilities, greatly facilitates the analysis when there are three or more traits, and is therefore well-suited to an empirical framework.

Third and last, in my model the socialization technology extends to both culturally homogamous and heterogamous households. This approach is not possible in the Bisin and Verdier model, which uses a unitary framework and for this reason assumes that heterogamous households have no available socialization technology (because then it would be unclear which culture the representative agent would want to transmit). Instead, in my model the technology of cultural capital production extends naturally to heterogamous households, yielding a trade-off between the socialization to the two parents' traits.

Regarding the production of the child's human capital, as in Chiappori et al. (2017) parents produce their child's human capital by using complementary time inputs. Furthermore, parental human capital improves the productivity of these time inputs. In their model, time investments into the child's human capital production are made at the expense of the household income – in my model, they are made at the expense of the child's cultural socialization.

1.4.3 Model analysis

With the technologies (1.15) and (1.16), the household problem (1.17)–(1.18) has closed-form solutions s_i^*, t_i^*. For the sake of clarity in the exposition, I make the following simplifying assumption.

Assumption 1: $\gamma_i = \dfrac{1}{T_i}$.

[15] This is because with only two traits, an increase in the population share of trait 1 mechanically leads to a decrease in the share of trait 2. Thus, a parent who cares about the population share of trait 1 must indirectly care about that of trait 2 as well. This is no longer true, however, when there are three or more traits. For instance, in the Bisin and Verdier framework Catholics could reduce their socialization effort if the population share of Protestants increases at the expense of Muslims'. This is not the case in my model.

This assumption imposes that a parent's time productivity in socialization or human capital formation reaches exactly 0 when they spend all of their time budget on only one activity. It guarantees interior solutions, while providing simpler formulas for the solutions s_i^* and t_i^*. I now describe these solutions as well as some of their properties, first for homogamous households, and then for heterogamous ones.

Homogamous households. In a homogamous household the two parents have aligned interests. They both wish to transmit their common culture as well as human capital to the child, although they may disagree on how much to favor one over the other. With n denoting the common trait of the two parents, the first-order conditions are

$$(\mu\,\nu_1 + \nu_2)\big(a_i(1-\gamma_i s_i^*) - a_0\,q_n\big) = (\mu+1)(b_i+h_i)(1-\gamma_i t_i^*) \qquad (i=1,2).$$

At the optimum, parent i's marginal returns from investing time in cultural capital or human capital formation should be equal. On the left-hand side, the marginal return from the socialization time s_i is increasing in the two parents' relative preferences for cultural capital ν_1 and ν_2, and in parent i's productivity a_i; and is decreasing in the productivity and intensity of oblique socialization, a_0 and q_n. On the right-hand side, the marginal return from time t_i spent on human capital formation is increasing in productivity b_i and the human capital of parent i, h_i.

The solution is obtained by using the saturated time constraint (1.18) and assumption 1:

$$s_i^* = T_i \times \frac{(\mu\,\nu_1+\nu_2)(a_i - a_0\,q_n)}{(\mu\,\nu_1+\nu_2)a_i + (\mu+1)(b_i+h_i)} \tag{1.19}$$

whenever this expression is positive (i.e. when $a_i > a_0 q_n$, so that parent i has an incentive to vertical socialization), and $s_i^* = 0$ otherwise. Note that the ratio in expression (1.19) is always inferior to 1, so that parent i never devotes their whole time budget to socialization. The following proposition describes how this optimal time allocation changes with the characteristics of the parents and of the population.

Proposition 1: In homogamous households, the time that parent i spends on cultural socialization is decreasing in his or her human capital level, h_i; and in the population share of the parents' common trait, q_n; and it is increasing in both parents' relative preference for cultural capital, ν_1 and ν_2. Furthermore, it is increasing in parent's 1

relative power μ if and only if $\nu_1 > \nu_2$, and decreasing otherwise.

The proof is obtained by differentiating the solution (1.19) with respect to the parameters of interest. Proposition 1 confirms that the model encapsulates the trade-offs between human capital formation and cultural socialization mentioned at the beginning of this section. Taking parental preferences as fixed, two types of time substitution occur in the model: parents with higher human capital reallocate their time toward human capital formation, as do cultural-majority parents. The first kind of substitution results from the comparative advantage of parents with higher human capital in the child's human capital production. The second is a consequence of vertical socialization coming at the expense of oblique socialization, which relates to the *cultural substitution* property introduced by Bisin and Verdier (2001).

If the preference for cultural capital changes for one parent, both parents respond to that change by reallocating their time in the same direction, as a consequence of the cooperativeness implied by the collective household framework. The same is true for changes in the power balance. Finally, the comparative statics with respect to the parental productivities a_i and b_i are straightforward: if their productivity in one activity increases, parents reallocate their time toward that activity.

Heterogamous households. In a heterogamous household, parents have different objectives: parent 1 wants to socialize the child to trait n_1 and parent 2 to trait n_2. Compared to the homogamous case, this divergence in the parents' interests modifies how they react to changes in the model's parameters. First, changes in parental preferences or in the power balance lead them to reallocate their time in opposite directions: when one parent increases their socialization time investment, the other parent reduces it. Second, a higher population share for any parent's trait decreases the incentive for vertical socialization for both parents, because vertical socialization happens at the expense of oblique socialization in all traits. Thus, if oblique socialization improves for one of the two parents, the household's incentive for vertical socialization decreases.

Formally, the first-order conditions are

$$\mu\,\nu_1\big(a_1(1-\gamma_1 s_1^*) - a_0\,q_{n_1}\big) - \nu_2\,a_0\,q_{n_2} = (\mu+1)(b_1+h_1)(1-\gamma_1 t_1^*)$$
$$\nu_2\big(a_2(1-\gamma_2 s_2^*) - a_0\,q_{n_2}\big) - \mu\,\nu_1\,a_0\,q_{n_1} = (\mu+1)(b_2+h_2)(1-\gamma_2 t_2^*).$$

These conditions are mostly similar to the homogamous case. The novelty is the adverse effect of vertical socialization from parent i on the oblique socialization to parent $-i$'s trait (in the first line for instance, this effect is represented by the term $-\nu_2\, a_0\, q_{n_2}$).

Once again, the solution is obtained by using the time constraint (1.18):

$$s_1^* = T_1 \times \frac{\mu\,\nu_1(a_1 - a_0\, q_{n_1}) - \nu_2\, a_0\, q_{n_2}}{\mu\,\nu_1\, a_1 + (\mu+1)(b_1+h_1)} \tag{1.20}$$

$$s_2^* = T_2 \times \frac{\nu_2(a_2 - a_0\, q_{n_2}) - \mu\,\nu_1\, a_0\, q_{n_1}}{\nu_2\, a_2 + (\mu+1)(b_2+h_2)}. \tag{1.21}$$

The following proposition describes the mechanisms at hand in heterogamous households: in addition to the two kinds of substitution occurring in homogamous households, parents must also adapt their time investments to suit their diverging interests.

Proposition 2: In heterogamous households, the time that parent i spends on cultural socialization is decreasing in his or her human capital level, h_i; in the population shares of the two parents' traits, q_{n_1} and q_{n_2}; and in the relative preference for cultural capital of the other parent, ν_{-i}; and it is increasing in his or her relative power, and in his or her relative preference for cultural capital ν_i.

Contrary to the homogamous case, a concession is made between the two parents in producing the child's human capital, which is a public good, versus producing cultural capital in the parents' respective traits, which is a private good enjoyed separately by each parent. The power balance notably determines the importance of socializing the child to parent 1's trait versus parent 2's trait. As a parent obtains more power, they dedicate more time to the cultural socialization of the child (their private good), while the other parent reallocates time toward human capital production (the public good).

This concludes the short analysis of the model. In Appendix 1.7.3 I discuss implications for household formation by using a matching framework, and for population dynamics.

1.5 Estimation

In this section I estimate the structural model developed in section 3.2. The method used is very similar to that of the reduced-form model from section 1.3. Indeed, I still exploit the variation in the traits of the respondents and their parents as a source of identification. This time however, I use not only the respondents' religious affiliation but also their educational attainment as an explained variable. This approach is possible because the structural model predicts both the religious socialization and the human capital of individuals. As I will explain in this section, the estimation framework can thus be understood as a mixture of a multinomial logit and an ordered logit model.

In section 1.5.1 I present the framework for the estimation. In section 1.5.2 I present the results.

1.5.1 Methodology

This section describes how I apply the model to the data. The religious affiliation of respondents in the data is explained by their predicted level of religious capital in the model, using a multinomial logit framework. Similarly, their educational attainment is explained by their predicted level of human capital in the model through an ordered logit framework. I then combine these two frameworks in the log-likelihood function for the estimation.

Measuring parents' human capital. Before delving into the estimation, a discrepancy between the model and the data needs to be addressed. In the model, the parents have a human capital trait h, which is continuous. In the data however, I measure this level of human capital by using the educational attainment variable, which is discrete. Thus, when solving the household program for two parents with observed educational attainments e_1 and e_2, I must decide how e_1 and e_2 translate into human capital levels h_1 and h_2.

As a simple solution, I assume that each educational attainment e is associated with a fixed human capital level $\tilde{h}_e$. Rather than choosing the $\tilde{h}_e$ exogenously however, I consider them as parameters to be estimated. In the model, any parent with educational attainment e is thus assumed to have the human capital level $\tilde{h}_e$.

Religious affiliation. For each individual i, the model predicts the cultural capital of i in any religion n, K_{in}, as a function of her parents' religious and educational traits and of the religions' population shares. To map these predictions onto the data, assume as in section 1.3 that i ultimately selects the religious affiliation

$$\arg\max_{n} \ \ln(K_{in}) + \varepsilon_{in}, \tag{1.22}$$

where the ε_{in} are distributed i.i.d. Gumbel. Again, the probability that i will select religion n is then given by

$$\pi_{in} = \frac{\exp(\ln K_{in})}{\sum_{\ell=1}^{N} \exp(\ln K_{i\ell})}. \tag{1.23}$$

This is a nonlinear multinomial logit model, in the sense that the probability π_{in} takes the standard softmax form, but the $\ln(K_{in})$ are nonlinear functions of the model's primitive parameters through the optimal time allocations (1.19)–(1.20)–(1.21).

Educational attainment. Similarly, for each individual i the model also predicts the level of human capital of i, H_i. Compared to the model of section 3.2, I add several parameters to describe this predicted level of human capital. These parameters intervene as additive constants in the log-human capital of children, and therefore they do not modify the optimal time allocation in the structural model: the solutions (1.19)–(1.20)–(1.21) remain unchanged. The first addition is simply a time trend parameter λ, which accounts for the fact that educational attainment has been rising over the period considered (c.f. Figure 1.1).

Second, I introduce parameters that reflect the baseline contributions of parents from different religious affiliations to the human capital formation of their children. These parameters are meant to capture religion-specific heterogeneity in how parents transmit human capital, which is unrelated to the specific trade-off between investments in human capital and cultural capital. The idea that religious ideology might influence human capital outcomes goes back at least to Weber's *The Protestant Ethic and the Spirit of Capitalism* (2013) [1905]. More recently, Botticini and Eckstein (2007) and Becker and Woessmann (2009) have provided evidence that the comparatively higher educational outcomes of Jews during the Middle-Ages and of Protestants during the late-19th century, respectively, could be attributed to a religious incentive to educate children. Becker et al.

(2020) also provided evidence that parents with a history of forced migration have more-educated children, which they explain by a stronger preference of such parents for mobile assets. In a historically Christian-majority country such as France, religious affiliation is correlated with migration ascendency, thus suggesting that patterns of investments in human capital could be dependent upon the parents' religion. In practice, I account for this heterogeneity by adding parental religion fixed effects, h_{1n} for mothers n and h_{2n} for fathers n, to the (log-)human capital of children. These fixed effects capture systematic differences in children's educational outcomes across parental religious affiliations while leaving space to identify trade-offs at the margin between investments in human versus religious capital as described in the model.

To map the predicted level of human capital onto the data, I consider H_i as the deterministic component of a latent variable which, in turn, predicts the educational level of i. Specifically, suppose that the actual level of (log-)human capital of individual i is $\ln(H_i) + \eta_i$, where η_i is a random shock. Suppose further that i attains the educational level e_i according to the rule

$$e_i = \begin{cases} 1 & \text{if} \quad \ln(H_i) + \eta_i \leq \bar{h}_1, \\ 2 & \text{if} \quad \bar{h}_1 < \ln(H_i) + \eta_i \leq \bar{h}_2, \\ \vdots & \\ E & \text{if} \quad \ln(H_i) + \eta_i > \bar{h}_{E-1}, \end{cases} \tag{1.24}$$

where E is the number of possible educational levels and $\bar{h}_1, \ldots, \bar{h}_{E-1}$ are parameters to be estimated. If the η_i are distributed i.i.d. logistic, this is an ordered logit model, such that the probability that i will attain the educational level e is given by

$$\phi_{ie} = \begin{cases} \dfrac{1}{1 + \exp(\ln(H_i) - \bar{h}_1)} & \text{if } e = 1, \\ \dfrac{1}{1 + \exp(\ln(H_i) - \bar{h}_2)} - \dfrac{1}{1 + \exp(\ln(H_i) - \bar{h}_1)} & \text{if } e = 2, \\ \vdots & \\ 1 - \dfrac{1}{1 + \exp(\ln(H_i) - \bar{h}_{E-1})} & \text{if } e = E. \end{cases} \tag{1.25}$$

Again, this model is nonlinear because H_i is a nonlinear function of the model's parameters.

Log-likelihood function and parametrization. Finally, suppose that the error terms ε_{in} and η_i are independent as well.[16] Then the probability that i will select the religious affiliation n and attain the educational level e is simply $\pi_{in} \times \phi_{ie}$. The model's log-likelihood is then

$$\ln L = \sum_i w_i \sum_{n=1}^{N} \sum_{e=1}^{E} \mathbf{1}_{\{i \text{ is } ne\}} \ln(\pi_{in} \times \phi_{ie}), \tag{1.26}$$

where the probabilities π_{in} and ϕ_{ie} implicitly depend on the model's parameters, and the w_i are sampling weights.

For the estimation, I impose two restrictions on the parametrization of the model presented in section 3.2. First, I assume that relative preferences for religious capital versus human capital are homogeneous within a gender–religion category. In other words, all mothers of religion n are assumed to have the same preference, denoted by ν_{1n}. Similarly, all fathers of religion n have the same preference ν_{2n}. This assumption is consistent with the model by Bisin and Verdier (2000), who assume that preferences are culture-specific constants. I extend their approach by supposing that within a culture, preferences may differ between men and women (which Bisin and Verdier could not do because they used a unitary household model). Second, I assume away the productivity difference between mothers and fathers, both in cultural socialization and human capital formation: $a_1 = a_2 = a$ and $b_1 = b_2 = b$. The reason for doing so is that differences in preferences between mothers and fathers are not precisely identified from differences in productivity between them. Indeed, as seen from the solutions to the household problem (1.19)–(1.20)–(1.21), identifying one from the other relies on the variation in the population shares q_{in}, which in practice does not seem sufficient to obtain robust estimates on different specifications. Thus a choice must be made to allow for gender heterogeneity in preferences or in productivity: here I choose the former.

I summarize the parametrization of the model under these additional assumptions, also considering assumption 1:

[16] In Appendix 1.7.4 I examine this hypothesis by analyzing the deviance residuals of the estimated model. I find that residuals on the religion and education dimensions are not strongly correlated, providing suggesting evidence that error terms are indeed independent.

$$\ln K_n = k_n + a\left(s_1 - \frac{s_1^2}{2T_1}\right)\mathbf{1}_{\{n_1=n\}} + a\left(s_2 - \frac{s_2^2}{2T_2}\right)\mathbf{1}_{\{n_2=n\}} + a_0\left(1 - s_1 - s_2\right)q_n$$

$$\ln H = \lambda(y - 1948) + h_{1n_1} + h_{2n_2} + (b + h_1)\left(t_1 - \frac{t_1^2}{2T_1}\right) + (b + h_2)\left(t_2 - \frac{t_2^2}{2T_2}\right)$$

$$u_i = \nu_{in_i}\ln(K_{n_i}) + \ln H.$$

For now, I exogenously fix the power balance in the couple by setting $\mu = 1$, so that the spouses have equal power. The parameters to estimate are thus the following:

- the relative preference for religious capital, ν_{1n} for mothers of religion n and ν_{2n} for fathers of religion n,
- the cultural adoption constants k_n for all n,
- the time productivities of religious socialization, a for vertical socialization by mothers and fathers, and a_0 for oblique socialization,
- the time productivity of human capital formation, b,
- the total time budgets of the parents, T_1 for mothers and T_2 for fathers,
- the human capital levels $\tilde{h}_e$ associated with the educational attainments e,
- the ordered logit thresholds $\bar{h}_e$,
- the religion-specific contributions to human capital, h_{1n} for mothers n and h_{2n} for fathers n,
- the time trend parameter in human capital formation, λ.

With $N = 5$ religions and $E = 3$ education levels, this makes a total of $5N + 2E + 5 = 36$ parameters. Of those, four are not identified. First, as in the reduced-form analysis of section 1.3, the k_n are identified only up to a common additive constant. Again, I normalize this constant to 0 for the most common denomination, Christians: $k_2 = 0$. Second, the time productivity in human capital formation of the lowest human capital level cannot be distinguished from the baseline time productivity of human capital formation.[17] As such, I normalize to 0 the added productivity of having a primary school diploma or less: $\tilde{h}_1 = 0$. Third, the religion-specific contributions to human capital are only identified up to a constant:[18] I also normalize those for Christian mothers and fathers, $h_{12} = h_{22} = 0$.

[17] The time allocation solutions depend on the sum $b_i + \tilde{h}_i$. Therefore, choosing the parameters $(b_1, b_2, \tilde{h}_e)$ or $(b_1 + \kappa, b_2 + \kappa, \tilde{h}_e - \kappa)$ leads to the same model outcomes.

[18] The parameters (h_{1n}, h_{2n}) and $(h_{1n} + \kappa, h_{2n} - \kappa)$ lead to the same human capital outcomes for any constant κ, so we need to anchor one of these parameters. Furthermore, the parameters $(h_{2n}, \bar{h}_e)$ and $(h_{2n} - \kappa, \bar{h}_e + \kappa)$ also lead to the same human capital outcomes, so one of them also needs to be fixed constant.

These normalizations leave 32 free parameters to estimate. Next, I compute the maximum likelihood estimator of these parameters with the log-likelihood expression (1.26). As in section 1.3, the covariance matrix is obtained via the BHHH estimator.

1.5.2 Results

Table 1.7 presents the estimation results. The fit can be compared to the null model with an intercept only for each religion–education type ($N \times E - 1 = 14$ free parameters), which has a deviance of 64 510. This comparison yields a pseudo-R^2 of 0.30.

I now turn to the estimated parameters. First, the estimates for the cultural adoption constants k_n are broadly consistent with the corresponding estimates in the reduced-form analysis. By default, individuals are most likely to select the No religion affiliation, followed by Christian, Other, Muslim and, finally, Jewish. As discussed in section 1.3, these results somewhat reflect the specificities of religious affiliations: for instance, while adopting the No religion trait requires little investment in religious capital, becoming Jewish without a Jewish parent is very rare.

Second, the relative values of religious capital for mothers and fathers, ν_{1n} and ν_{2n}, exhibit wide differences across gender and religions. Overall, the estimates suggest that mothers have stronger preferences for religion versus education than fathers do, except among Nones. They also suggest that Muslims and Jews value religious capital the most relative to education-oriented human capital, followed by Others, Nones, and finally Christians. Note that the estimates for Jewish mothers and fathers are very large but remain very imprecisely estimated. These large standard errors are a consequence of data being scarce for Jewish parents, particularly in the context of heterogamous households, which are used to identify the difference in preferences between mothers and fathers.

Third, the estimates of productivity in religious socialization, a and a_0, suggest that vertical socialization operates on a larger order of magnitude than oblique socialization, confirming the important roles of parents in the socialization process.

Moving to the estimates related to human capital formation, the measures of added productivity from parental human capital $\tilde{h}_e$ are, reassuringly, increasing in the associated educational attainment e. The added productivity obtained from holding a Secondary diploma is approximately the same as that from holding a Tertiary diploma. This result confirms that higher-educated parents are more productive when spending time to trans-

Table 1.7: Structural model of religious socialization and human capital formation, Estimates

	Estimates
	model (1.26)
Religious socialization	
Cultural adoption constants k_n, by religion n	
None	2.08 (0.13)
Christian	0 (baseline)
Muslim	−1.91 (0.22)
Jewish	−3.19 (0.32)
Other	−1.29 (0.21)
Value of religious capital for mothers ν_{1n}, by religion n	
None	0.08 (0.04)
Christian	0.68 (0.10)
Muslim	2.13 (0.46)
Jewish	36.22 (55.60)
Other	1.59 (0.33)
Value of religious capital for fathers ν_{2n}, by religion n	
None	0.14 (0.05)
Christian	0.09 (0.05)
Muslim	1.98 (0.35)
Jewish	17.07 (10.15)
Other	0.00 (0.24)
Vertical religious socialization time productivity, a	11.74 (2.90)
Oblique religious socialization time productivity, a_0	1.60 (0.33)
Human capital formation	
Contributions to human capital from mothers h_{1n}, by religion n	
None	−0.45 (0.08)
Christian	0 (baseline)
Muslim	1.85 (0.28)
Jewish	4.82 (1.19)
Other	1.16 (0.19)
Contributions to human capital from fathers h_{2n}, by religion n	
None	−0.48 (0.07)
Christian	0 (baseline)
Muslim	0.88 (0.22)
Jewish	5.01 (0.74)
Other	−0.81 (0.16)
Human capital formation time productivity, b	14.68 (3.99)
Added productivity $\tilde{h}_e$, by education level e	
Primary	0 (baseline)
Secondary	1.85 (0.44)
Tertiary	3.75 (0.88)
Human capital threshold $\bar{h}_1$: Primary → Secondary	6.90 (0.95)
Human capital threshold $\bar{h}_2$: Secondary → Tertiary	11.07 (0.95)
Other	
Total time budget of mothers, T_1	0.66 (0.15)
Total time budget of fathers, T_2	0.68 (0.16)
Time trend, λ	0.01 (0.00)
Observations	18 155
Sampling weights	Yes
Deviance ($-2 \ln L$)	45 247
Pseudo-R^2	0.30

Note: Standard errors in parentheses.

mit human capital to their children. This feature of the estimates is, of course, driven by the fact that higher-educated parents have higher-educated children in the data. The baseline time productivity in human capital formation, b, is similar in magnitude but greater than the time productivity of vertical socialization, a.

Regarding the religion-specific parental contributions to human capital, I find that relative to Christians, Jewish mothers and fathers contribute the most to the human capital formation of children, confirming the documented strong emphasis on human capital among Jews (Botticini and Eckstein 2007, 2012). Muslims also contribute more than Christians by default, and therefore, their lower educational rates must be attributable to the trade-off between investments in religion versus education. None parents contribute less than Christians. Finally, among Other parents, Other mothers contribute more than Christian ones, while Other fathers contribute less.

1.5.3 Interpretation of the estimates

With Cobb–Douglas preferences, the parameter ν_i can be interpreted in this way: parent i is indifferent between the child's religious capital K_{n_i} increasing by 1%, and the child's human capital H increasing by ν_i%. However, since both H and K_{n_i} are latent variables with no obvious measure scale, this interpretation is not immediately helpful in understanding parents' trade-offs in terms of religious transmission and education. Instead, another way to understand this trade-off is to ask a question such as: what loss in their children's educational attainment are parents accepting in exchange for a 1% increase in the chance that they will transmit their religious affiliation? With this question we finally tackle the motivating question of the paper, that is, the cost that parents pay to transmit their religion.

Note that a parent's transmission probability (the probability that the child will share his or her religious affiliation) depends not only on that parent's gender and religion but also on their education level, on the spouse's characteristics, and on the religions' population shares which are relevant for the child's socialization. All these factors will play a role in determining how costly religious transmission is for a given parent. To keep the illustration simple, I therefore focus on parents in homogamous households, and I take as a baseline the population shares for each religion in the year 2008, corresponding to the year of the survey. Furthermore, I consider households in which both parents share

Parents' education	Parents' religion					
	None	Christian	Muslim	Jewish	Other	Musl./Chri. ratio
	Marginal rate of transformation					
Primary or less	0.16	0.25	2.09	–	0.30	8.4
Secondary	0.33	0.41	3.91	–	0.64	9.5
Tertiary or more	0.38	0.35	4.59	–	0.78	13.1

Table 1.8: Cost of religious transmission, in terms of child's probability of Tertiary education.

the same education level. Finally, since preference parameters for Jewish parents are very imprecisely estimated, I do not include them in this analysis.

Figure 1.18 shows, for homogamous households in which both parents have a Secondary education, what this trade-off looks like. The cost of religious transmission is measured in terms of foregone probability that the child obtains a Tertiary education, through a marginal rate of transformation. At their predicted transmission profile (the white dots on the graph), Christian parents, for instance, are renouncing a 0.41 percentage point (p.p.) chance that their child will attain a Tertiary education, in order to increase by 1 p.p. the chance that they will transmit their religious affiliation. This number is to 0.33 for Nones, rising to 0.64 for Others, and to 3.91 for Muslims. Thus, Muslims are paying a marginal price approximately 10 times greater than Christians in terms of their child's educational attainment to ensure the transmission of their religious affiliation.

In Table 1.8 I summarize this information on marginal rates of transformation for different levels of parental education. As suggested by the evidence on Secondary-educated parents, I find that Muslim parents of all educational levels face higher costs than other denominations overall. Keeping the example of Muslims and Christians for illustration, the ratio of their marginal rates of transformation varies between 8 and 13. This strongly suggests that Muslims do indeed pay a significantly larger price than other denominations to transmit their religious affiliation. The difference between other affiliations is not as striking, although I find that Others pay a steeper cost than Christians, who themselves have a cost comparable to Nones.

1.5.4 Log-likelihood decomposition

Even though the estimated parameters from Table 1.7 provide a rough idea of the magnitude of the different mechanisms at play in the model, a more detailed analysis is needed

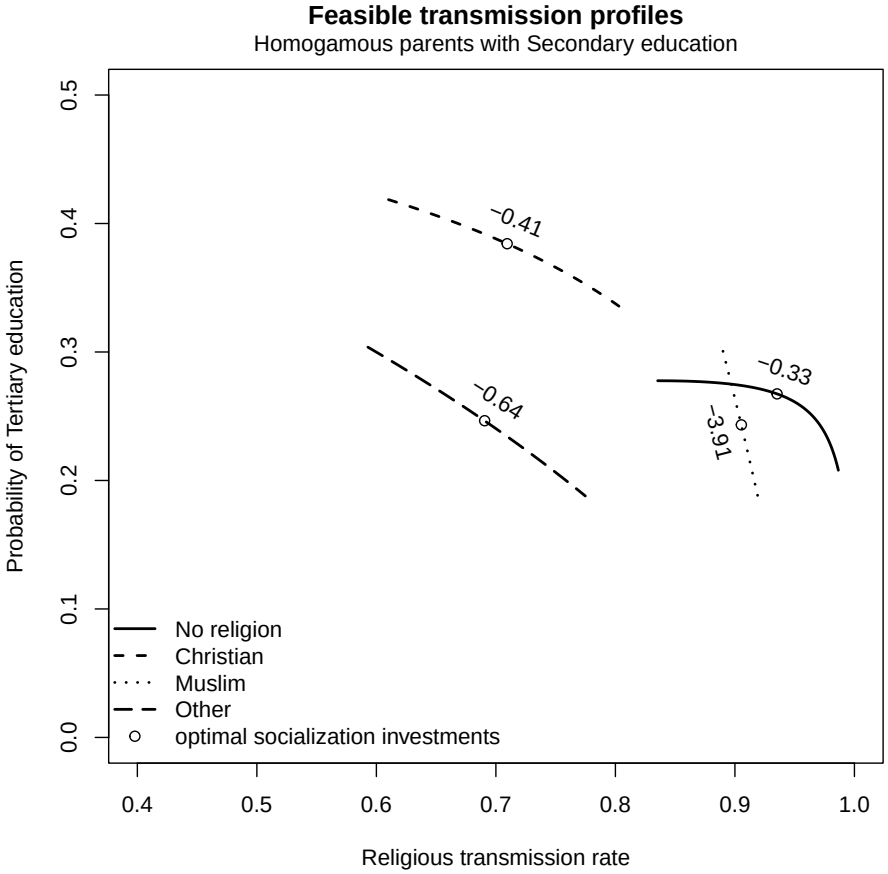

Figure 1.18: Trade-off between the religious transmission rate and the probability of the child attaining a Tertiary education. The lines depict the feasible transmission profiles (technology frontier) for households with Secondary-educated parents in religiously homogamous households. The white dots correspond to the profiles chosen by parents as predicted by the model. The marginal rate of transformation at this chosen profile is reported for each household type. (I do not report the corresponding results for Jewish households given the imprecision of the estimates.)

to understand their respective importance. Here I delve deeper into this issue with a log-likelihood decomposition, which allows me to rank the three mechanisms at play in the model by order of importance in terms of explanatory power. These three mechanisms are (1) heterogeneous parental preferences across gender and religious affiliations, (2) the role of parental human capital in the substitution effect, and (3) oblique socialization. In order to measure the respective explanatory power of these three mechanisms, I consider three corresponding restrictions on the model parameters. Such restrictions allow me to shut down the three mechanisms separately and, as a result, to measure their respective ability to explain the variation in the data. The mechanisms and associated restrictions are the following.

Parental preferences. What if parents had homogeneous preferences? I restrict all preference parameters to take the same value, $\nu_{1n} = \nu_{2n} = \nu$ for all n and for some constant ν. For ν, I choose the value of which is associated with the maximum log-likelihood under this constraint and while keeping other parameter estimates constant.

Parental human capital. What if parental human capital did not affect time productivity in the child's human capital formation? In this case I set the added time productivity due to parental human capital, $\tilde{h}_1$ and $\tilde{h}_2$, to 0.

Oblique socialization. What if oblique socialization played no role in the transmission process? In order to cancel this effect I set the oblique socialization parameter a_0 to 0.

I then use the log-likelihood from equation (1.26), which I denote here using the lowercase ℓ, to obtain a measure of explanatory power for each of these mechanisms. Specifically, I consider as a baseline the difference between the log-likelihood evaluated at the actual estimator $\hat{\boldsymbol{\beta}}$ (the one reported in Table 1.7), and the log-likelihood of the null model: $\ell(\hat{\boldsymbol{\beta}}) - \ell_0$.[19] We saw that the structural model performs better than the null model at explaining the data, such that this difference is positive. Then, for each mechanism I modify the vector of estimated parameters $\hat{\boldsymbol{\beta}}$ by applying the associated parameter restriction, yielding a new vector of parameters $\hat{\boldsymbol{\beta}}_{\text{restr}}$. I can then compute the difference $\ell(\hat{\boldsymbol{\beta}}_{\text{restr}}) - \ell_0$ and compare it to the baseline difference $\ell(\hat{\boldsymbol{\beta}}) - \ell_0$ (this baseline is necessarily greater by definition of the maximum likelihood estimator). If $\ell(\hat{\boldsymbol{\beta}}_{\text{restr}}) - \ell_0$ is close to $\ell(\hat{\boldsymbol{\beta}}) - \ell_0$ it means that the mechanism which was shut down has little explanatory power. Conversely, if the two are far from each other, it means that the mechanism which was shut down actually matters a lot to explain variation in the data. Having noted this, a possible statistic to measure the explanatory power of the different mechanisms is therefore

$$1 - \frac{\ell(\hat{\boldsymbol{\beta}}_{\text{restr}}) - \ell_0}{\ell(\hat{\boldsymbol{\beta}}) - \ell_0}. \tag{1.27}$$

A higher value of this statistic means a larger explanatory power for the associated mechanism.

The values of this statistic for the three mechanisms considered are presented in Figure

[19] Recall that the null model is defined by an intercept only for each religion–education type of the child.

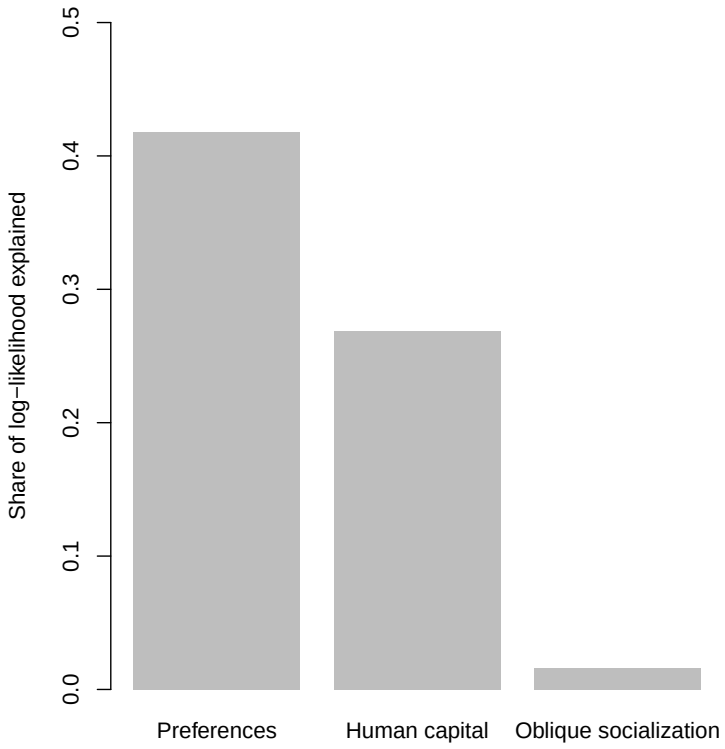

Figure 1.19: Three mechanisms ranked in terms of explanatory power.

1.19. I find that the dominant mechanism in terms of explanatory power is parental preferences, with a decrease of 42% in log-likelihood when shutting it down. Parental human capital comes second, with a decrease of 27% in log-likelihood when ignoring its effect. Finally, oblique socialization seems to play a minimal role, with a decrease of only 2% in log-likelihood when shutting it down. This may however be due in part to the rough measure of oblique socialization used in the model, namely population shares at the national level.

1.6 Conclusion

While there is historical evidence of investments in religious transmission to children coming at the expense of their human capital outcomes, research has not really truly addressed this issue in a contemporary context. This gap remains despite strong anecdotical evidence that such trade-offs take place, especially among religious minorities – in the US for instance, among the Amish, Jehovah's Witnesses, or Hasidic Jews. In this paper, I have used data on religious affiliation in contemporary France to study this trade-off. I first documented how parents from different affiliations contribute to religious transmission to their children, finding that mothers transmit more than fathers, religious minorities transmit more than majorities, and higher-educated parents transmit less than lower-educated parents. To explain these stylized facts, I built a structural model that illustrates the trade-offs between investments in religious versus human capital. The estimates of the structural model suggest, for instance, that Muslims pay a cost that can be more than 10 times greater than that paid by Christians to transmit their religion, in terms of the educational attainment of their children.

More work remains to understand how this trade-off occurs across different contexts, such as other countries, different religious affiliations not well represented in the French context, and even for other cultural traits such as language or ethnicity. Additional data on religion and how parents allocate their time could also could also help refine the estimates obtained from the new methodology developed here. In particular, the model would benefit greatly from local measures of religions' population shares to better understand the religious environment in which individuals grow up; intensive measures of individuals' religion, such as the intensity of their religious practices, to better gauge their involvement with their declared religious affiliation; or measures of parental investments in the culture or formal education of their children to better understand how parents substitute their investments in various dimensions. Overall, this work lays out an interesting research program for better understanding the costs that parents pay to transmit their culture.

Bibliography

Abdelgadir, A. and Fouka, V. (2020), 'Political secularism and Muslim integration in the West: Assessing the effects of the French Headscarf Ban', *American Political Science Review* **114**(3), 707–723.

Abdelhadi, E. (2019), 'The hijab and muslim women's employment in the united states', *Research in Social Stratification and Mobility* **61**, 26–37.

Abramitzky, R., Boustan, L. and Eriksson, K. (2020), 'Do immigrants assimilate more slowly today than in the past?', *American Economic Review: Insights* **2**(1), 125–41.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2010), 'Identifying barriers to muslim integration in france', *Proceedings of the National Academy of Sciences* **107**(52), 22384–22390.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2016), *Why Muslim integration fails in Christian-heritage societies*, Harvard University Press.

Aizer, A. and Cunha, F. (2012), The production of human capital: Endowments, investments and fertility. NBER Working Paper 18429.

Ajbli, F. (2011), Les Françaises musulmanes face à l'emploi: le cas des pratiquantes "voilées" dans la métropole lilloise, PhD thesis, Paris, EHESS.

Akerlof, G. A. and Kranton, R. E. (2000), 'Economics and identity', *The quarterly journal of economics* **115**(3), 715–753.

Aksoy, O. (2017), 'Motherhood, sex of the offspring, and religious signaling', *Sociological Science* **4**, 511–527.

Aksoy, O. and Gambetta, D. (2016), 'Behind the veil: the strategic use of religious garb', *European Sociological Review* **32**(6), 792–806.

Aksoy, O. and Gambetta, D. (2021), 'The politics behind the veil', *European Sociological Review* **37**(1), 67–88.

Algan, Y., Malgouyres, C., Mayer, T. and Thoenig, M. (2022), 'The economic incentives of cultural transmission: Spatial evidence from naming patterns across france', *The Economic Journal* **132**(642), 437–470.

Andriantsimbazovina, J., Kabou, P. et al. (2020), *Laïcité et défense de l'État de droit*, Presses de l'Université Toulouse 1 Capitole.

Atasoy, Y. (2006), 'Governing women's morality: a study of islamic veiling in Canada', *European Journal of Cultural Studies* **9**(2), 203–221.

Atkin, D., Colson-Sihra, E. and Shayo, M. (2021), 'How do we choose our identity? a revealed preference approach using food consumption', *Journal of Political Economy* **129**(4), 1193–1251.

Beauchemin, C., Hamel, C. and Simon, P. (2016), *Trajectoires et origines: enquête sur la diversité des populations en France*, INED éditions.

Becker, G. S. (1965), 'A theory of the allocation of time', *The Economic Journal* **75**(299), 493–517.

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

Becker, G. S. (1974), 'A theory of marriage: Part II', *Journal of Political Economy* **82**(2), S11–S26.

Becker, S. O., Grosfeld, I., Grosjean, P., Voigtländer, N. and Zhuravskaya, E. (2020), 'Forced migration and human capital: Evidence from post-WWII population transfers', *The American Economic Review* **110**(5), 1430–63.

Becker, S. O. and Woessmann, L. (2009), 'Was Weber wrong? A human capital theory of Protestant economic history', *The Quarterly Journal of Economics* **124**(2), 531–596.

Bellemare, M. F., Novak, L. and Steinmetz, T. L. (2015), 'All in the family: Explaining the persistence of female genital cutting in west africa', *Journal of Development Economics* **116**, 252–265.

Benjamini, Y. and Hochberg, Y. (1995), 'Controlling the false discovery rate: a practical and powerful approach to multiple testing', *Journal of the Royal statistical society: series B (Methodological)* **57**(1), 289–300.

Benzer, T. (2022), Removing cultural barriers to education: State-run islamic schools and girls' education in turkey, Working paper.

Berndt, E., Hall, B., Hall, R. and Hausman, J. (1974), 'Estimation and inference in nonlinear structural models', *Annals of Economic and Social Measurement* **3**(4), 653–665.

Bessin, M. (2002), *Autopsie du service militaire 1965–2001*, Autrement.

Biavaschi, C., Giulietti, C. and Siddique, Z. (2017), 'The economic payoff of name americanization', *Journal of Labor Economics* **35**(4), 1089–1116.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2011), 'Ethnic identity and labour market outcomes of immigrants in europe', *Economic Policy* **26**(65), 57–92.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2016), 'Bend it like beckham: Ethnic identity and integration', *European Economic Review* **90**, 146–164.

Bisin, A. and Topa, G. (2003), 'Empirical models of cultural transmission', *Journal of the European Economic Association* **1**(2/3), 363–375.

Bisin, A., Topa, G. and Verdier, T. (2004), 'Religious intermarriage and socialization in the United States', *Journal of Political Economy* **112**(3), 615–664.

Bisin, A. and Tura, G. (2022), Marriage, fertility, and cultural integration in Italy. Working Paper.

Bisin, A. and Verdier, T. (2000), '"Beyond the Melting Pot": Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits', *The Quarterly Journal of Economics* **115**(3), 955–988.

Bisin, A. and Verdier, T. (2001), 'The economics of cultural transmission and the dynamics of preferences', *Journal of Economic Theory* **97**, 298–319.

Bisin, A. and Verdier, T. (2011), The economics of cultural transmission and socialization, *in* 'Handbook of Social Economics', Vol. 1, North-Holland, pp. 339–416.

Bisin, A. and Verdier, T. (2022), Advances in the economic theory of cultural transmission. NBER Working Paper 30466.

Bolt, U., French, E., Maccuish, J. H. and O'Dea, C. (2021), The intergenerational elasticity of earnings: exploring the mechanisms, Cepr discussion paper no. dp15975.

Botticini, M. and Eckstein, Z. (2007), 'From farmers to merchants, conversion and diaspora: Human capital and Jewish history', *Journal of the European Economic Association* **5**(5), 885–926.

Botticini, M. and Eckstein, Z. (2012), *The Chosen Few*, Princeton University Press.

Bourdieu, P. (1979), 'Les trois états du capital culturel', *Actes de la recherche en sciences sociales* **30**(1), 3–6.

Bouzar, D. and Bouzar, L. (2009), *Allah a-t-il sa place dans l'entreprise?*, Albin Michel.

Carvalho, J.-P. (2013), 'Veiling', *The Quarterly Journal of Economics* **128**(1), 337–370.

Carvalho, J.-P., Koyama, M. and Sacks, M. (2017), 'Education, identity, and community: lessons from the Jewish emancipation', *Public Choice* **171**, 119–143.

Carvalho, J.-P., Koyama, M. and Williams, C. (2022), Resisting education. Oxford Discussion paper 982.

Carvalho, J.-P. and McBride, M. (2022), The formation of religious beliefs and preferences, *in* K. Zimmerman, ed., 'Handbook of Labor, Human Resources and Population Economics', Springer.

Chaudhary, L. and Rubin, J. (2011), 'Reading, writing, and religion: Institutions and human capital formation', *Journal of Comparative Economics* **39**, 17–33.

Cheung, M.-W. and Wu, J. (2018), 'On the probabilistic transmission of continuous cultural traits', *Journal of Economic Theory* **174**, 300–323.

Chiappori, P.-A. (1992), 'Collective labor supply and welfare', *Journal of Political Economy* **100**(3), 437–467.

Chiappori, P.-A. (2017), *Matching with transfers: The economics of love and marriage*, Princeton University Press.

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

Choo, E. and Siow, A. (2006), 'Who marries whom and why', *Journal of Political Economy* **11'**(1), 175–201.

Cintas, C., Gosse, B. and Vatteville, É. (2012), Quand l'identité religieuse devient une préoccupation du management des ressources humaines, *in* 'Management et religion: Décryptage d'un lien indéfectible', EMS Editions, pp. 83–98.

Corekcioglu, G. (2021), 'Unveiling the effects of a headscarf ban: Evidence from municipal jobs in turkey', *Journal of Comparative Economics* **49**(2), 382–404.

Cunha, F. and Heckman, J. (2007), 'The technology of skill formation', *The American Economic Review* **97**(2), 31–47.

Del Boca, D., Flinn, C. and Wiswall, M. (2014), 'Household choices and child development', *The Review of Economic Studies* **81**(1), 137–185.

Del Boca, D., Flinn, C. and Wiswall, M. (2016), 'Transfers to households with children and child development', *The Economic Journal* **126**(596), F136–F183.

Del Bono, E., Francesconi, M., Kelly, Y. and Sacker, A. (2016), 'Early maternal time investments and early child outcomes', *The Economic Journal* **126**, F96–F135.

Droogsma, R. A. (2007), 'Redefining hijab: American muslim women's standpoints on veiling', *Journal of Applied Communication Research* **35**(3), 294–319.

Drydakis, N. (2013), 'The effect of ethnic identity on the employment of immigrants', *Review of Economics of the Household* **11**(2), 285–308.

Fan, X. and Wu, L. (2022), The shaping of a gender norm: Marriage, labor, and foot-binding in historical china, Working paper.

Fernández-Reino, M., Di Stasio, V. and Veit, S. (2022), 'Discrimination unveiled: A field experiment on the barriers faced by muslim women in germany, the netherlands, and spain', *European Sociological Review* .

Fouka, V. (2020), 'Backlash: The unintended effects of language prohibition in U.S. schools after World War I', *The Review of Economic Studies* **87**(1), 204–239.

Fryer, R. G. and Levitt, S. D. (2004), 'The causes and consequences of distinctively black names', *The Quarterly Journal of Economics* **119**(3), 767–805.

Galichon, A., Kominers, S. D. and Weber, S. (2019), 'Costly concessions: An empirical framework for matching with imperfectly transferable utility', *Journal of Political Economy* **127**(6), 2875–2925.

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

Garre, F. G., Vermunt, J. K. and Croon, M. A. (2002), 'Likelihood-ratio tests for order-restricted log-linear models: A comparison of asymptotic and bootstrap methods', *Metodologa de las Ciencias del Comportamiento* **4**(1), 1–18.

Gaspard, F. and Khosrokhavar, F. (1995), *Le foulard et la République*, FeniXX.

Gibbons, R. and Waldman, M. (2004), 'Task-specific human capital', *The American Economic Review* **94**(2), 203–207.

Greene, W. H. (2008), *Econometric analysis*, 6th edn, Pearson Education, Inc.

Gresh, A. (2020), 'Islam: A one-sided debate'.
URL: *https://orientxxi.info/dossiers-et-series/islam-a-one-sided-debate,4281*

Guiso, L., Sapienza, P. and Zingales, L. (2006), 'Does culture affect economic outcomes?', *Journal of Economic Perspectives* **20**(2), 23–48.

Gulesci, S., Jindani, S., La Ferrara, E., Smerdon, D., Sulaiman, M. and Young, H. (2021), A stepping stone approach to understanding harmful norms, CEPR Discussion Paper No. DP15776.

Harrison, K. A. (2016), 'Hiding under the veil of" dress policy": Muslim women, hijab, and employment discrimination in the united states', *Geo. J. Gender & L.* **17**, 831.

Heckman, J., Pinto, R. and Savelyev, P. (2013), 'Understanding the mechanisms through which an influential early childhood program boosted adult outcomes', *American Economic Review* **103**(6), 2052–86.

Hu, Z. and Wang, Z. (2021), Nutrition, labor supply, and productivity: Evidence from ramadan in indonesia, Working paper.

Iannaccone, L. R. (1990), 'Religious practice: A human capital approach', *Journal for the Scientific Study of Religion* **29**(3), 297–314.

Iannaccone, L. R. (1992), 'Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives', *Journal of Political Economy* **100**(2), 271–291.

Iannaccone, L. R. (1998), 'Introduction to the economics of religion', *Journal of Economic Literature* **36**(3), 1465–1495.

Institut Français d'Opinion Publique [IFOP] (2019), Etude auprès de la population musulmane en france, 30 ans après l'affaire des foulards de creil, Technical report.

Institut Montaigne (2014–2021), Religion au travail: croire au dialogue. baromètre du fait religieux en entreprise 2014–2021, Technical report.

Institut Montaigne (2016), A french islam is possible, Technical report.

Iyer, S. (2016), 'The new economics of religion', *Journal of Economic Literature* **54**(2), 395–441.

Jia, R. and Persson, T. (2021), 'Choosing ethnicity: The interplay between individual and social motives', *Journal of the European Economic Association* **19**(2), 1203–1248.

Jouili, J. S. (2020), *Pious Practice and Secular Constraints*, Stanford University Press.

Katz, D., Baptista, J., Azen, S. and Pike, M. (1978), 'Obtaining confidence intervals for the risk ratio in cohort studies', *Biometrics* **34**(3), 469–474.

Khosrokhavar, F. (2004), L'islam des jeunes filles en france, *in* 'Le foulard islamique en questions', Éditions Amsterdam, pp. 89–94.

Kirkebøen, L., Leuven, E. and Mogstad, M. (2022), 'College as a marriage market', *NBER Working Paper* (28688).

La Boussole (2004), *Annuaire des mosquées de France*.

Lecomte, J.-P. (2001), Représentations et réalités des fonctions sociales du service militaire dans la société française (1868–2001). Thèse de doctorat en Science politique soutenue en 2001 à Paris, Institut d'études politiques.

Livny, A. (2020), *Trust and the Islamic advantage: religious-based movements in Turkey and the Muslim world*, Cambridge University Press.

Maillard, D. (2017), *Quand la religion s'invite dans l'entreprise*, Fayard.

Maurin, E. and Navarrete-Hernandez, N. (2023), 'Behind the veil: the effect of banning the islamic veil in schools', *Economic Policy* **38**(113), 63–98.

McFadden, D. (1973), Conditional logit analysis of qualitative choice behavior, *in* P. Zarembka, ed., 'Frontiers in Econometrics', Academic Press, New York, chapter 4, pp. 105–142.

Montgomery, J. D. (2010), 'Intergenerational cultural transmission as an evolutionary game', *American Economic Journal: Microeconomics* **2**(4), 115–136.

Morsy, M. (1989), *Les femmes du Prophète*, FeniXX.

Munshi, K. (2011), 'Strength in numbers: Networks as a solution to occupational traps', *The Review of Economic Studies* **78**(3), 1069–1101.

Munshi, K. (2019), 'Caste and the Indian economy', *Journal of Economic Literature* **57**(4), 781–834.

Nordmann, C. (2004), *Le foulard islamique en questions*, Éditions Amsterdam.

Novak, L. (2020), 'Persistent norms and tipping points: The case of female genital cutting', *Journal of Economic Behavior & Organization* **177**, 433–474.

Patacchini, E. and Zenou, Y. (2016), 'Social networks and parental behavior in the intergenerational transmission of religion', *Quantitative Economics* **7**, 969–995.

Patel, D. S. (2012), 'Concealing to reveal: The informational role of islamic dress', *Rationality and Society* **24**(3), 295–323.

Platteau, J.-P. (2017), *Islam instrumentalized*, Cambridge University Press.

Read, J. G. and Bartkowski, J. P. (2000), 'To veil or not to veil? a case study of identity negotiation among muslim women in austin, texas', *Gender & society* **14**(3), 395–417.

Roth, A. E. and Sotomayor, M. A. O. (1990), *Two-sided matching: A study in game-theoretic modeling and analysis*, Cambridge University Press.

Sakalli, S. O. (2019), Secularization and religious backlash: Evidence from turkey, Technical report, Technical Report, Working Paper.

Saleh, M. (2016), 'Public mass modern education, religion, and human capital in twentieth-century Egypt', *The Journal of Economic History* **76**(3), 697–735.

Saleh, M. (2018), 'On the road to heaven: Taxation, conversions, and the Coptic–Muslim socioeconomic gap in Medieval Egypt', *The Journal of Economic History* **78**(2), 394–434.

Scott, J. W. (2009), *The politics of the veil*, Vol. 7, Princeton University Press.

Sénat (2021), 'Compte rendu intégral: Séance du mardi 30 mars 2021', *Journal Officiel de la République française* .

Shofia, N. M. (2020), Why veil? religious headscarves and the public role of women, Working paper.

Silhouette-Dercourt, V., Sy, O. S. and Desjeux, D. (2019), 'Cosmopolitan veiling in paris: Young french muslim women in transition', *Youth and Globalization* **1**(1), 65–87.

Siow, A. (2015), 'Testing Becker's theory of positive assortative matching', *Journal of Labor Economics* **33**(2), 409–441.

Skrondal, A. and Rabe-Hesketh, S. (2009), 'Prediction in multilevel generalized linear models', *Journal of the Royal Statistical Society Series A: Statistics in Society* **172**(3), 659–687.

Squicciarini, M. (2020), 'Devotion and development: Religiosity, education, and economic progress in 19th-century France', *The American Economic Review* **110**(11), 3454–3491.

Valfort, M.-A. (2020), 'Anti-muslim discrimination in france: Evidence from a field experiment', *World Development* **135**, 105022.

van der Hasselt, G. (2019), 'The muslim veil in france: Why so controversial? three questions to hakim el kaouri'.
URL: *https://www.institutmontaigne.org/en/blog/muslim-veil-france-why-so-controversial*

Weber, M. (2013), *The Protestant Ethic and the Spirit of Capitalism*, Routledge.

Weichselbaumer, D. (2020), 'Multiple discrimination against female immigrants wearing headscarves', *ILR Review* **73**(3), 600–627.

1.7 Appendix

1.7.1 Descriptive statistics for the TeO survey

General descriptives

Table 1.1: General descriptive statistics of the TeO survey.

	Mean	StdDev	Min	Max	Obs.
Age	36	11.5	17	60	21,761
Female (%)	52.8				21,761
Education (%)					21,761
Primary or less	8.0				–
Secondary	63.6				–
Tertiary or more	28.4				–
Religion (%)					21,443
No religion	29.3				–
Christian	39.2				–
Muslim	26.6				–
Jewish	0.8				–
Other	4.1				–
Partner					
Has partner (%)	72.5				21,761
Same-sex partner[1] (%)	0.7				13,242
Raised by... (%, several may apply)					21,761
Both parents	86.1				
Mother only	14.9				
Father only	2.3				
Mother's education (%)					20,239
Primary or less	59.3				–
Secondary	30.4				–
Tertiary or more	10.2				–
Father's education (%)					19,239
Primary or less	54.2				–
Secondary	31.2				–
Tertiary or more	14.7				–
Parents' religion					
Homogamous parents (same religion, %)	89.3				20,671
Shares religion with at least one parent (%)	84.9				20,988

Notes: [1] Information only available if the partner lives in the same house.

Education

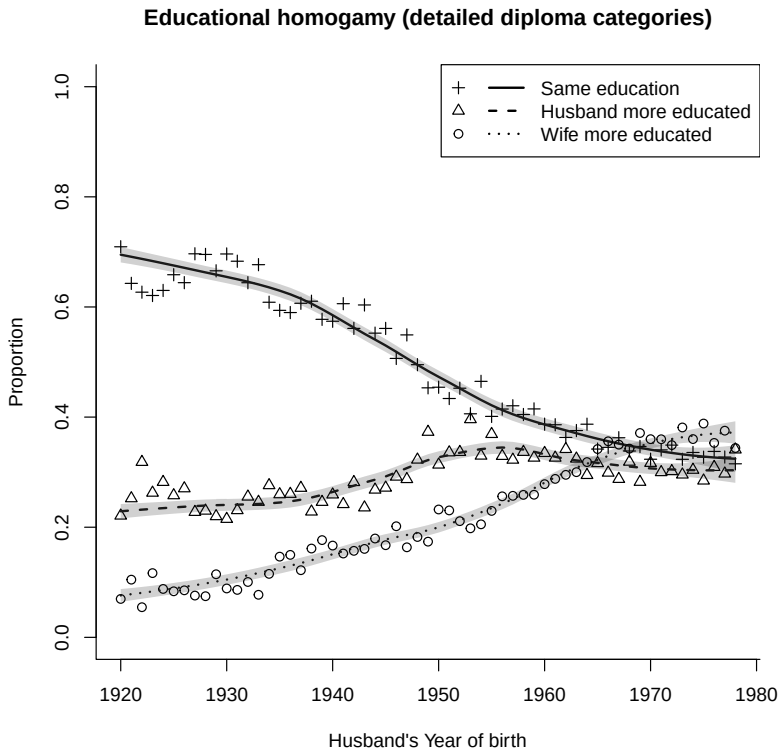

Figure 1.1: Educational homogamy with detailed diploma categories.

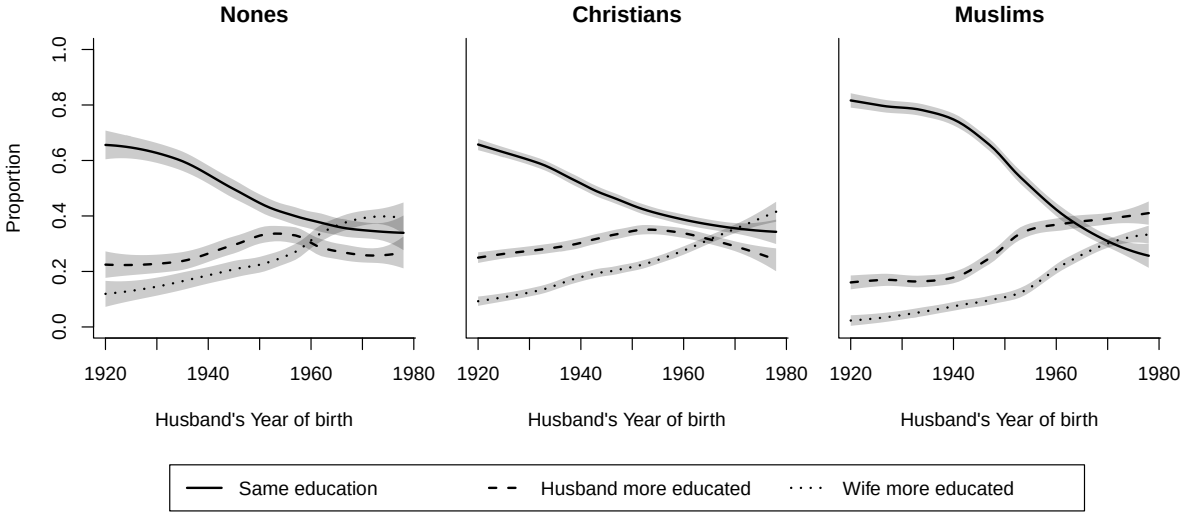

Figure 1.2: Educational homogamy with detailed diploma categories, same-religion couples.

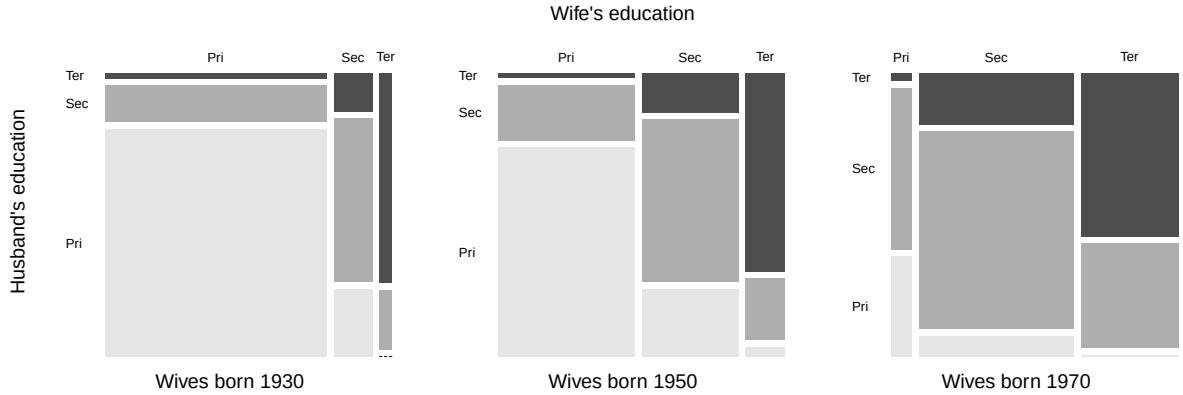

Figure 1.3: Educational assortment in couples with a wife born in 1930, 1950, and 1970.

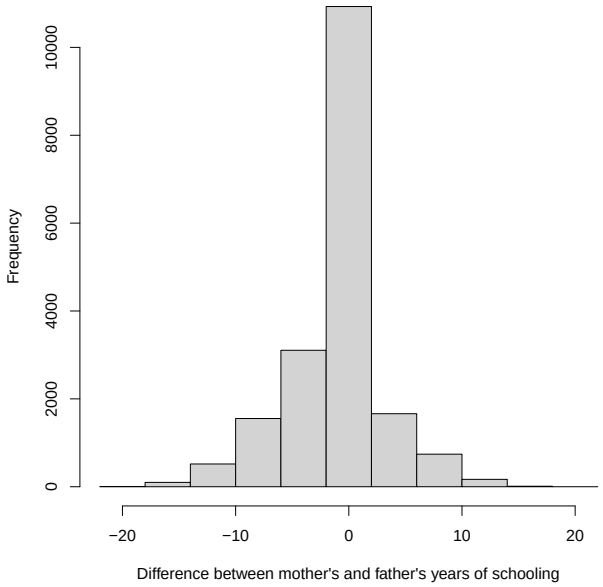

Figure 1.4: Distribution of the difference in years of schooling between mothers and fathers.

Religion

See Figures 1.5, 1.6, 1.7; and Table 1.2.

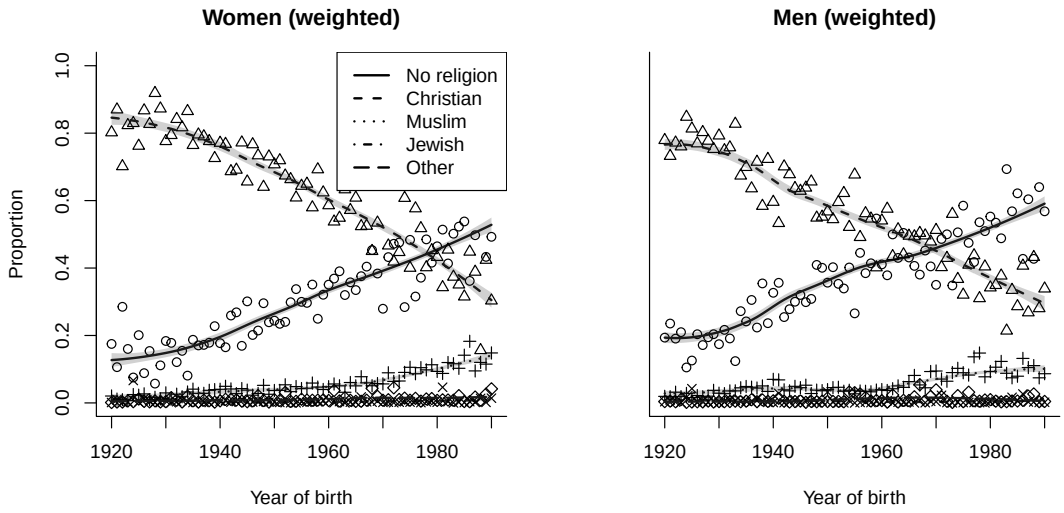

Figure 1.5: Religious affiliation, Women and Men (using sampling weights).

Table 1.2: Religious affiliations and homogamy.

Mother's religion	Father's religion					Total	Homogamy
	None	Christian	Muslim	Jewish	Other		
None	2448	221	105	8	28	2810	0.87
Christian	1071	9044	240	32	89	10476	0.86
Muslim	118	42	5905	0	3	6068	0.97
Jewish	9	25	4	149	1	188	0.79
Other	110	76	19	1	923	1129	0.82
Total	3756	9408	6273	190	1044	20671	
Homogamy	0.65	0.96	0.94	0.78	0.88		

Note: For each line, homogamy is computed as the ratio of mothers in a homogamous union divided by the total number of mothers in that line (idem for fathers in each column). Homogamy rates can thus differ within a single religion between mothers and fathers because of they have different distributions regarding religion.

Figure 1.6: Religious assortment in couples with a wife born in 1930, 1950, and 1970 (using sampling weights).

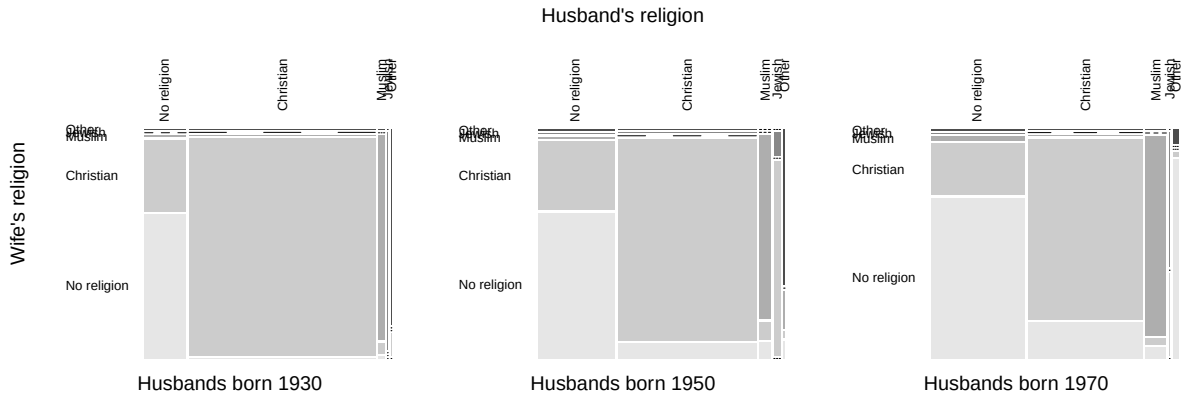

Figure 1.7: Religious assortment in couples with a husband born in 1930, 1950, and 1970 (using sampling weights).

Transmission of education

See Tables 1.3 and 1.4 for additional regressions.

Let's further investigate the interaction between the mother's and father's education levels, e_m and e_f, in determining the education of the child e_c. To do this, let's simplify the education variable even more than above by defining e_i as

$$e_i = \mathbf{1}_{\{i \text{ has (at least) a Secondary diploma}\}}.$$

Call $\mu_{h_m h_f} = \mathbb{P}(e_c = 1 \mid e_m, e_f)$ the probability that a child has (at least) a Secondary diploma, conditional on her mother and father having education levels e_m and e_f respectively. A simple measure of the interaction effect between the parents' education levels is then

$$\mu_{11} - \mu_{10} - \mu_{01} + \mu_{00}. \tag{1.28}$$

Table 1.3: Transmission of education (Ordered Logit).

	Child's education			
	(Ord. logit)	(Ord. logit)	(Ord. logit)	(Ord. logit)
Mother's education				
Secondary	0.64 (0.02)	1.04 (0.03)	0.97 (0.03)	0.98 (0.03)
Tertiary	1.00 (0.03)	2.60 (0.10)	2.57 (0.10)	2.57 (0.10)
Father's education				
Secondary	0.63 (0.02)	0.99 (0.03)	0.96 (0.03)	0.96 (0.03)
Tertiary	1.56 (0.03)	1.74 (0.05)	1.72 (0.05)	1.72 (0.05)
Mother's × Father's education				
Secondary × Secondary		−0.76 (0.04)	−0.68 (0.04)	−0.67 (0.04)
Secondary × Tertiary		−0.40 (0.06)	−0.36 (0.07)	−0.36 (0.07)
Tertiary × Secondary		−2.03 (0.11)	−2.10 (0.11)	−2.10 (0.11)
Tertiary × Tertiary		−1.76 (0.12)	−1.75 (0.12)	−1.76 (0.12)
Mother's religion				
Christian			0.25 (0.02)	0.24 (0.03)
Muslim			0.03 (0.28)	0.02 (0.49)
Jewish			−0.02 (0.16)	1.47 (1.18)
Other			0.63 (0.16)	0.59 (0.43)
Father's religion				
Christian			0.28 (0.02)	0.28 (0.05)
Muslim			−0.11 (0.28)	0.03 (0.66)
Jewish			1.23 (0.16)	1.05 (0.31)
Other			−0.74 (0.21)	−1.03 (1.49)
Mother's × Father's religion				
Christian × Christian				0.00 (0.06)
Christian × Muslim				−0.19 (0.78)
Christian × Jewish				0.35 (0.43)
Christian × Other				1.98 (2.04)
Muslim × Christian				−0.13 (1.41)
Muslim × Muslim				−0.13 (0.82)
Muslim × Jewish				no data
Muslim × Other				1.52 (15.83)
Jewish × Christian				−1.76 (1.21)
Jewish × Muslim				−1.36 (5.92)
Jewish × Jewish				−1.35 (1.23)
Jewish × Other				no data
Other × Christian				0.55 (0.48)
Other × Muslim				−1.15 (2.00)
Other × Jewish				no data
Other × Other				0.21 (1.55)
Child's year of birth /100	0.30 (0.06)	0.40 (0.06)	1.07 (0.07)	1.08 (0.07)
Cut-off: Primary → Secondary	3.56 (1.24)	5.58 (1.25)	19.03 (1.28)	19.10 (1.28)
Cut-off: Secondary → Tertiary	7.70 (1.24)	9.76 (1.25)	23.25 (1.28)	23.33 (1.28)
Observations	18 793	18 793	18 222	18 222
Sampling weights	Yes	Yes	Yes	Yes
Residual Deviance	27098	26947	25901	25888

Note: Standard errors in parentheses.
Reference category for mother/father education is "Primary."
Reference category for mother/father religion is "No religion."

Table 1.4: Transmission of education (OLS).

	Child's education		
	(OLS)	(OLS)	(OLS)
Mother's education	0.13***	0.21***	0.20***
	(0.01)	(0.02)	(0.02)
Father's education	0.18***	0.25***	0.24***
	(0.01)	(0.01)	(0.02)
Mother's × Father's education		−0.04***	−0.04***
		(0.01)	(0.01)
Mother's religion...			
Christian			0.06***
			(0.01)
Muslim			0.01
			(0.05)
Jewish			−0.01
			(0.07)
Other			0.15**
			(0.05)
Father's religion...			
Christian			0.06***
			(0.01)
Muslim			−0.02
			(0.04)
Jewish			0.30***
			(0.07)
Other			−0.18**
			(0.06)
Child's year of birth /100	0.11**	0.08*	0.23***
	(0.03)	(0.03)	(0.03)
Observations	18793	18793	18222
Sampling weights	Yes	Yes	Yes
Adjusted R^2	0.14	0.15	0.16

Note: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Reference category for wife/husband religion fixed effects is "No religion."

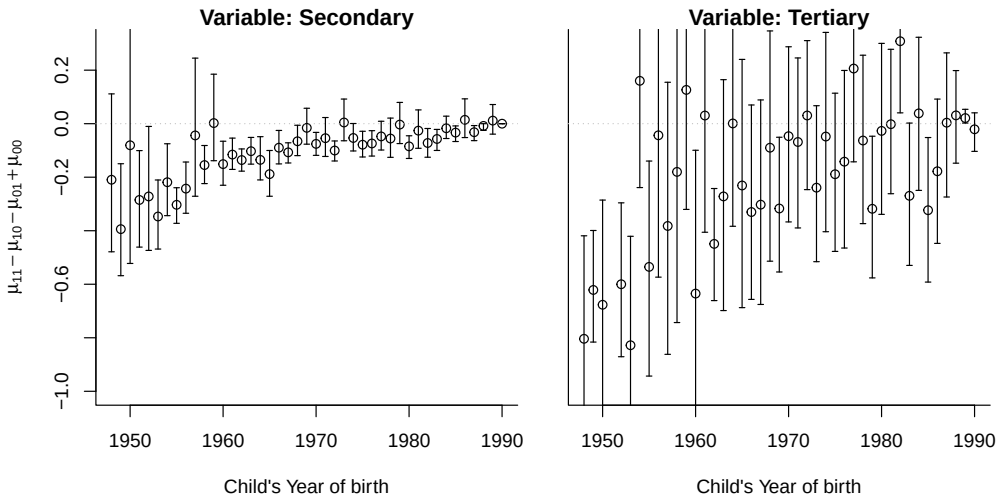

Figure 1.8: Interaction effects of parents' education levels for the child's education. 95% confidence intervals for $\mu_{11} - \mu_{10} - \mu_{01} + \mu_{00}$ are reported for each cohort. Left panel uses Secondary diplomas to define the binary education variable e_i, right panel uses Tertiary diplomas.

(For instance, in the linear probability model $\mu_{h_m h_f} = \alpha + \beta_m e_m + \beta_f e_f + \gamma e_m e_f$, we have $\mu_{11} - \mu_{10} - \mu_{01} + \mu_{00} = \gamma$.) I estimate the expression (1.28) on the whole sample first. The estimator $\hat{\mu}_{e_m e_f}$ of $\mu_{e_m e_f}$ is the sample mean of e_c on the subsample of respondents with a mother e_m and a father e_f. The point estimate for (1.28) is then simply $\hat{\mu}_{11} - \hat{\mu}_{10} - \hat{\mu}_{01} + \hat{\mu}_{00}$. The confidence interval is obtained by simulation, knowing that each $\hat{\mu}_{e_m e_f}$ follows a binomial distribution. I obtain the point estimate -0.120, with $[-0.132, -0.108]$ for the 95% confidence interval. This estimate can be interpreted as follows: the gain from having an additional Secondary-educated parent is 12 p.p. less for children who already have one Secondary-educated parent, compared to children who have none. This result indicates that interaction effects are negative. Next I perform the same exercise within cohorts. The results are shown in Figure 1.8. Again, estimates for (1.28) are negative, even within cohorts.

As a last control, I perform the same exercise but instead define e_i as

$$e_i = \mathbf{1}_{\{i \text{ has a Tertiary diploma}\}}.$$

Estimation of (1.28) on the full sample yields the point estimate -0.122 with 95% confidence interval $[-0.178, -0.066]$. Estimation within cohorts is again reported in Figure 1.8. Most point estimates remain negative, although many cannot be statistically distinguished from 0.

Migration

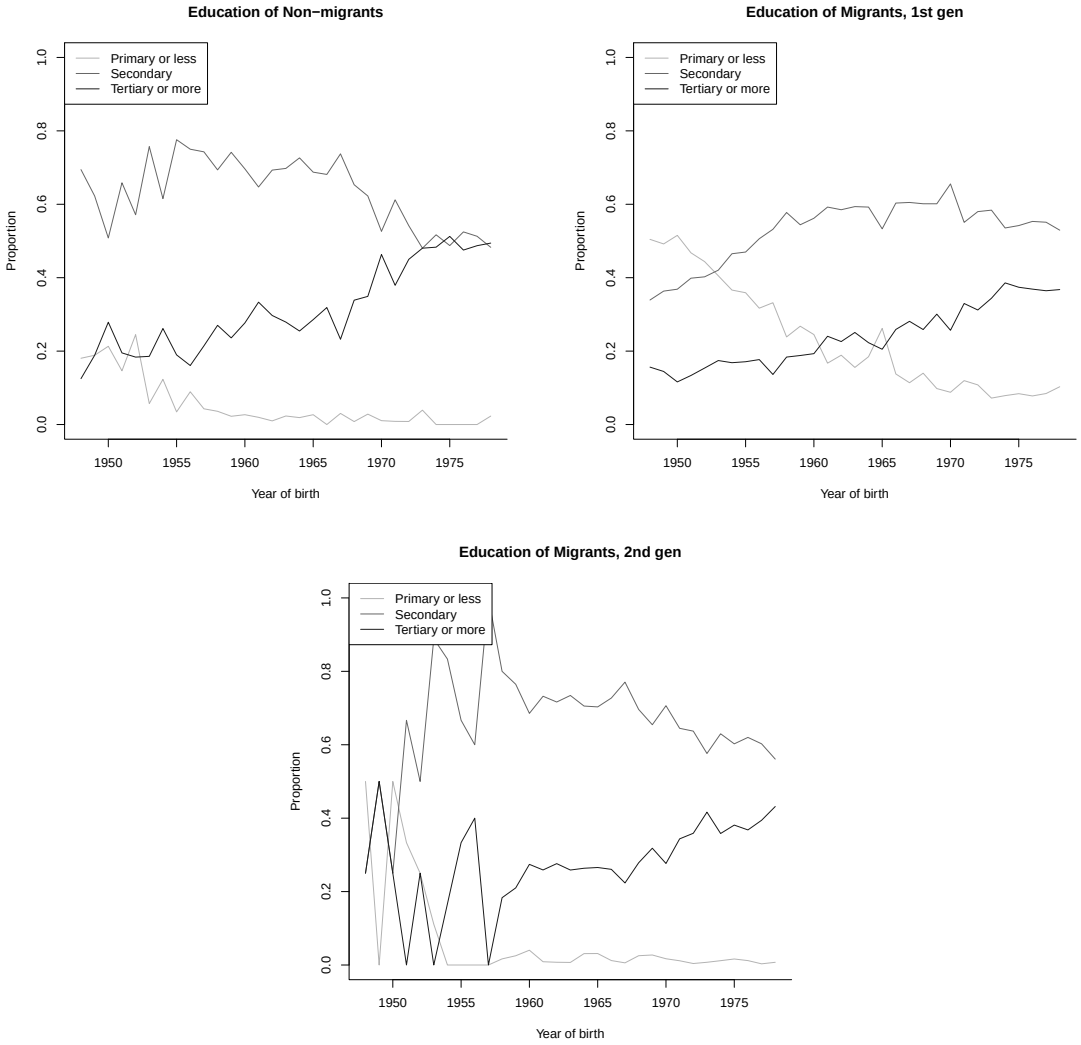

Figure 1.9: Education by Migration status.

Transmission of religion

Homogamy advantage. When focusing on households without a None parent, homogamous households perform significantly better than heterogamous households in passing on religious traits (Figure 1.10).

This advantage is also confirmed when considering transmission rates for any combination of parental religious affiliations (Figure 1.11).

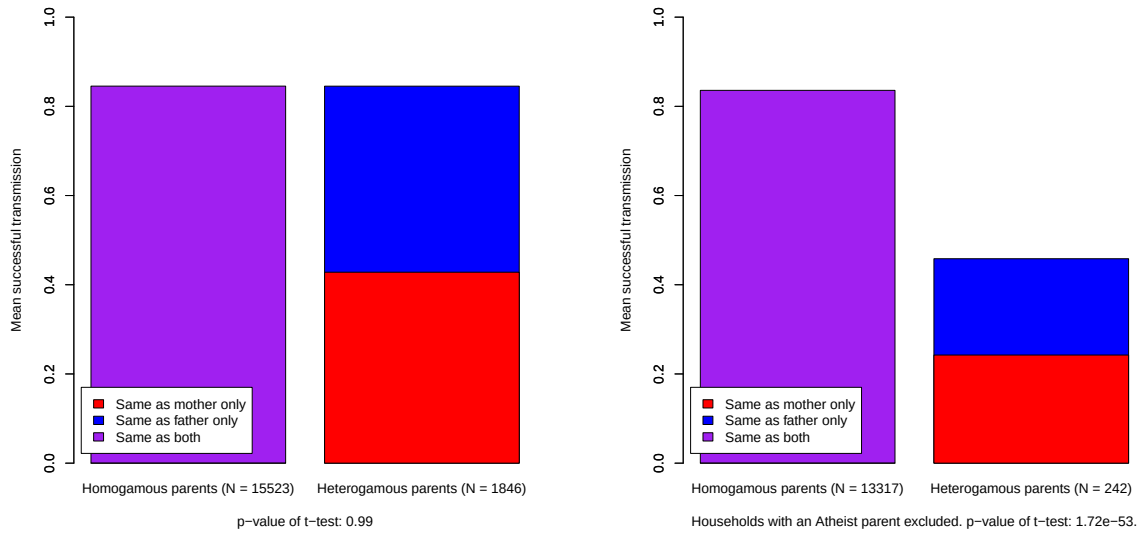

Figure 1.10: Transmission rates for homogamous and heterogamous households. The right-hand graph omits the respondents who declared having a 'No religion' parent.

Gender asymmetry. Another documented fact is that mothers pass on their cultural trait at a higher rate than fathers do. This difference is somewhat visible in Figure 1.11, where mothers' transmission rates (in red) seem overall more prominent than fathers' (in blue). However, no clear pattern emerges from the aggregated evidence. This is because the distribution of religious traits is different for mothers and fathers in the sample: in particular, there are more None fathers than mothers, which biases transmission success in the favor of fathers given the trend towards No religion mentioned above. For this reason, we must examine how mothers and fathers perform when they are in comparable situations. I systematically investigate this mother–father asymmetry in Figure 1.12 by comparing the respective religious transmission rates of mothers and fathers in symmetric household configurations. Specifically, for any religious traits a and b, I compute the difference between the transmission rate of mothers in households ab (i.e. when the mother has religion a and the father religion b) and the transmission rate of fathers in households ba (i.e. when the father has religion a and the mother religion b). I find that an argument can

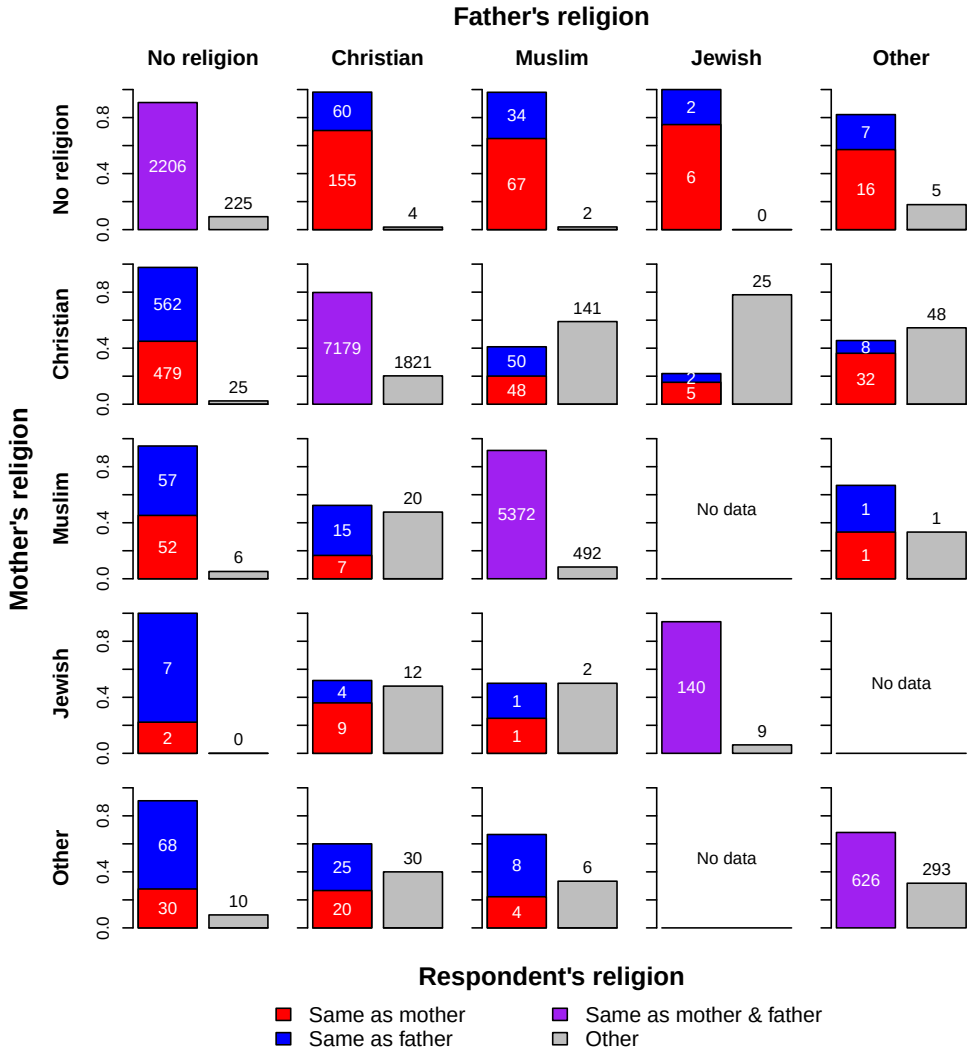

Figure 1.11: Transmission rates for all combinations of the parents' religions (number of observations reported for each bar).

indeed be made for the larger role of mothers in religious transmission: in five cases this difference in transmission rates is significant at the 95% level in favor of mothers (None vs. Christian, None vs. Muslim, Christian vs. None, Jewish vs. Christian, Other vs. Christian). The Jewish vs. Christian case is notable, as it reflects that Jewish affiliation is passed down from the mother and not the father. In contrast, there is no significant advantage for fathers at the 95% level. If we broaden the confidence interval to the 90% level, mothers gain a significant advantage in the Muslim vs. None case, while fathers gain a significant advantage in the Christian vs. Muslim case (perhaps reflecting the fact that Muslim affiliation is primarily passed down from the father).

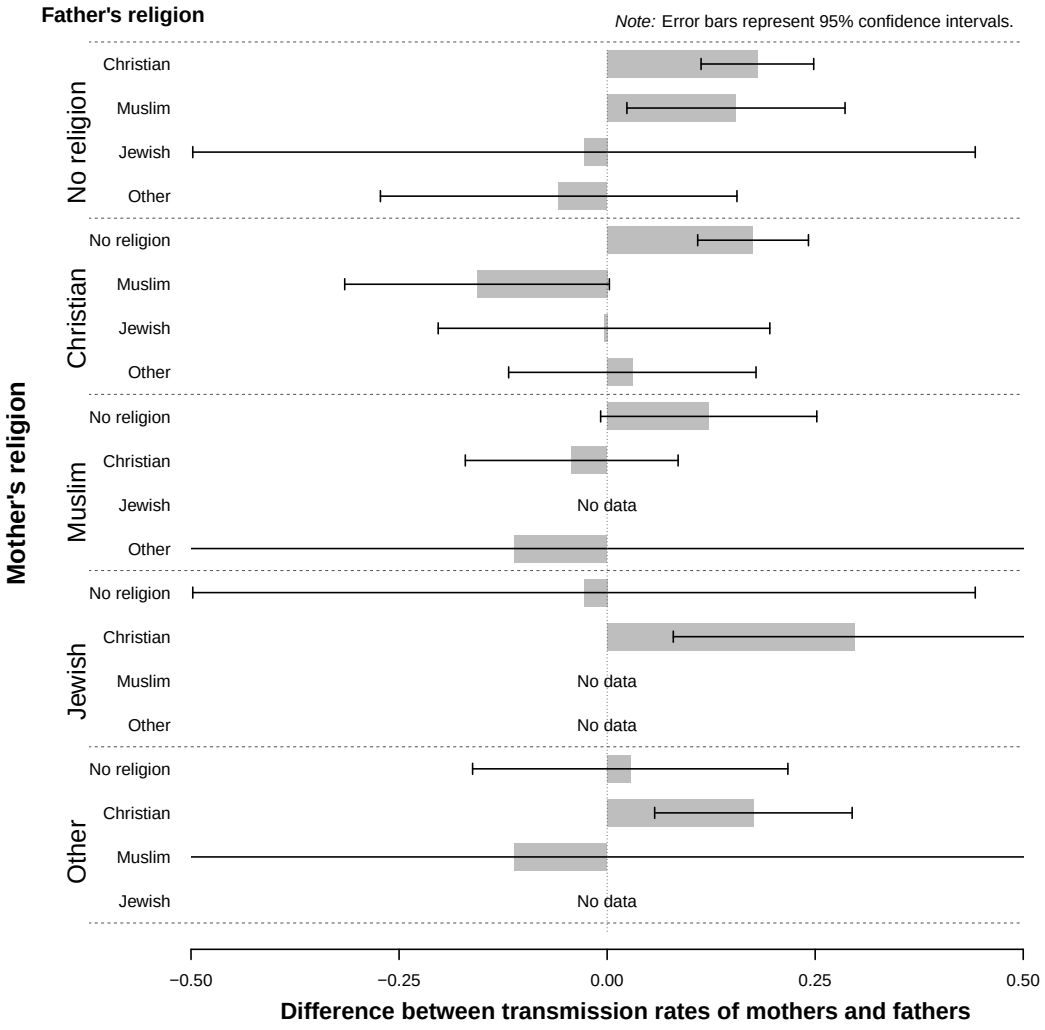

Figure 1.12: Mothers' advantage in religious transmission. Plotted are the differences between the transmission rate of mothers in households ab and the transmission rate of fathers in households ba for any religions $a \neq b$.

1.7.2 Reduced-form models of transmission

Reconstructing population shares

Finding which population shares to use is not straightforward. Ideally, one should use a time series of religious shares in France over the period considered. Unfortunately, this information is not consistently available for every year. Instead, I resort to using the TeO survey data to reconstruct these population shares. I assume that for a given birth cohort y, the population that contributes to oblique socialization for that cohort consists of all individuals born between $y-1$ and $y-60$ who were residents of metropolitan France no later than $y+18$. Population shares for each trait n are computed accordingly in that subsample (which includes respondents' parents) by using sampling weights. The limit of $y-60$ means that individuals born more than 60 years before a given cohort do not affect that cohort's oblique socialization. This limit is chosen somewhat arbitrarily to account for deaths among older individuals, given that dates of death are not available. Furthermore, behind the decision to count only residents at $y+18$ is the implicit assumption that religious affiliation is decided by age 18. The resulting population shares involved in oblique socialization for every birth cohort are shown in Figure 1.15. As a point of comparison, in the same figure I also show the corresponding estimates from the 2005 World Values Survey based on 996 respondents. (The 2005 population shares correspond to those involved in the oblique socialization of cohort $2005-18=1992$). The estimates for the shares of Nones and Christians differ, with approximately 12 p.p. more Nones in the World Values Survey than in the reconstructed shares. However, these shares are consistent with estimates from other studies.

Derivation of testable restrictions

Following equation (1.3), the independence of irrelevant alternatives assumption of the conditional logit model takes the following form:

$$\ln\left(\frac{\pi_{in}}{\pi_{i\ell}}\right) = \ln K_{in} - \ln K_{i\ell} = k_n + m_n \mathbf{1}_{\{i\text{'s mother is } n\}} + f_n \mathbf{1}_{\{i\text{'s father is } n\}} + \alpha\, q_{y_i n} - k_\ell - m_\ell \mathbf{1}_{\{i\text{'s mother is } \ell\}} - f_\ell \mathbf{1}_{\{i\text{'s father is } \ell\}} - \alpha\, q_{y_i \ell}$$

where I have made explicit that q_{in} depends on i only through her year of birth y_i.

Call $\pi_{in\,|\,yab}$ the probability that i adopts trait n conditional on belonging to birth cohort y, and having a mother a and a father b. Then for any two traits a and b and two birth cohorts y

and $\tilde{y}$, we have:

$$\ln\left(\frac{\pi_{ia\,|\,yaa}}{\pi_{ib\,|\,yaa}}\right) = k_a + m_a + f_a + \alpha\, q_{ya} - k_b - \alpha\, q_{yb} \tag{1.29}$$

$$\ln\left(\frac{\pi_{ia\,|\,yab}}{\pi_{ib\,|\,yab}}\right) = k_a + m_a + \alpha\, q_{ya} - k_b - f_b - \alpha\, q_{yb} \tag{1.30}$$

$$\ln\left(\frac{\pi_{ia\,|\,\tilde{y}ba}}{\pi_{ib\,|\,\tilde{y}ba}}\right) = k_a + f_a + \alpha\, q_{\tilde{y}a} - k_b - m_b - \alpha\, q_{\tilde{y}b} \tag{1.31}$$

$$\ln\left(\frac{\pi_{ia\,|\,\tilde{y}bb}}{\pi_{ib\,|\,\tilde{y}bb}}\right) = k_a + \alpha\, q_{\tilde{y}a} - k_b - m_b - f_b - \alpha\, q_{\tilde{y}b}. \tag{1.32}$$

It follows that

$$\ln\left(\frac{\pi_{ia\,|\,yaa}}{\pi_{ib\,|\,yaa}}\right) - \ln\left(\frac{\pi_{ia\,|\,yab}}{\pi_{ib\,|\,yab}}\right) - \ln\left(\frac{\pi_{ia\,|\,\tilde{y}ba}}{\pi_{ib\,|\,\tilde{y}ba}}\right) + \ln\left(\frac{\pi_{ia\,|\,\tilde{y}bb}}{\pi_{ib\,|\,\tilde{y}bb}}\right) = 0. \tag{1.33}$$

Note that we cannot take a reference trait n_0 as pivot, in the sense that if equation (1.33) is true for the traits an_0 and bn_0, it does not imply that it is true for the traits ab. This is because this equation involves different subpopulations depending on the choice of the two traits:

- if we consider the property (1.33) for the traits a and n_0, then the subpopulation involved consists of all individuals with parents aa, an_0, n_0a, or n_0n_0;
- for the traits b and n_0, it is the individuals with parents bb, bn_0, n_0b, or n_0n_0;
- for the traits a and b, it is the individuals with parents aa, ab, ba, or bb.

Since the first two points do not involve individuals with parents ab or ba, there is no way that any combination of the two associated equations would yield results on this subpopulation and, consequently, no way that they could imply (1.33) for traits a and b.

We can however take a birth cohort y_0 as a pivot. That is, equation (1.33) is true for all a, b, y, and $\tilde{y}$, if and only if it is true for all a, b, and y, but taking $\tilde{y} = y_0$ fixed. In practice, however, this approach is not useful as I do not have enough observations to perform the test for every cohort. Instead, I consider the approximate test

$$\ln\left(\frac{\pi_{ia\,|\,aa}}{\pi_{ib\,|\,aa}}\right) - \ln\left(\frac{\pi_{ia\,|\,ab}}{\pi_{ib\,|\,ab}}\right) - \ln\left(\frac{\pi_{ia\,|\,ba}}{\pi_{ib\,|\,ba}}\right) + \ln\left(\frac{\pi_{ia\,|\,bb}}{\pi_{ib\,|\,bb}}\right) = 0 \tag{1.34}$$

where $\pi_{in\,|\,ab}$ is the probability that i will adopt n conditional on having a mother a and a father b (but no longer conditioning on birth cohorts). This simplification relies on almost-constant population shares over the period considered. Test results are presented in Figure 1.1, with 100,000 parametric bootstrap simulations to obtain confidence intervals. The $\pi_{in\,|\,ab}$ are computed considering sampling weights. There are 10 tests to perform, 2 of which cannot be computed because of a lack of observations. Among 8 computable tests, 5 do not reject the

restriction at the 5% level, and 3 do.

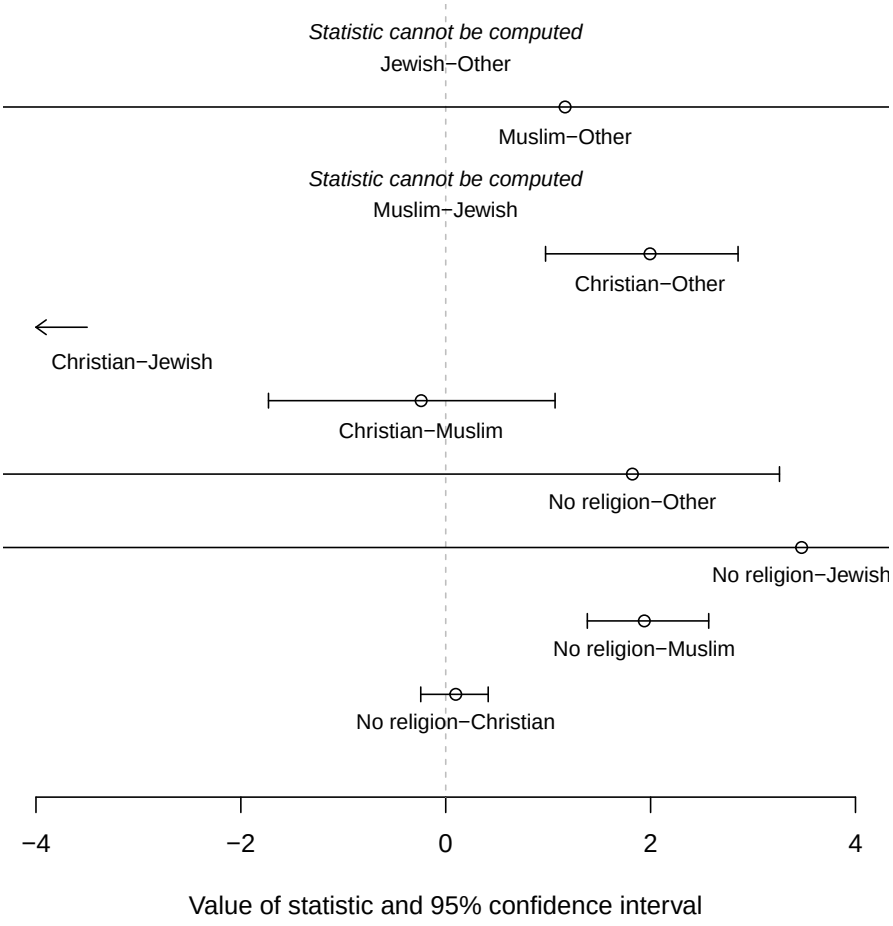

Figure 1.1: Statistical tests of equation (1.34) for all trait combinations ab.

As a next step, I test these hypotheses simultaneously through both global and multiple testing procedures. First, I compute the individual p-value for each test corresponding to a trait combination ab. This computation is not straightforward since it is not *a priori* clear which distribution the left-hand term of equation (1.34) follows under the null. Here I rely on a result from Katz et al. (1978), who show that a log-ratio of binomial distributions is approximately normally distributed. Since the $\pi_{in\,|\,ab}$ follow binomial distributions, each logarithm term in (1.34) is approximately normally distributed, and therefore, the left-hand side of equation (1.34) is approximately normally distributed (as a sum of normal distributions). I compute the p-value for the test ab by checking how the empirical estimate of the statistic compares to the normal distribution with mean zero and standard deviation equal to that recovered by parametric bootstrap (although it is not clear whether the null (1.34) would also modify the standard

deviation of the distribution). There are 8 such computable p-values (corresponding to the 8 computable tests).

Once p-values are computed, I can follow Bonferroni's method for global testing. Among 8 tests, 3 have a p-value below .05/8 (the 3 tests that reject the null individually), so Bonferroni's method leads to rejection (1.34). I can also follow the procedure of Benjamini and Hochberg (1995) for multiple testing, to control for the false discovery rate (FDR). In Figure 1.2 I plot the p-values corresponding to each test, ordered from smallest to largest, along with the threshold line of slope $j\gamma/S$, with S the total number of tests, j the index variable, and $\gamma = 5\%$ the level of the multiple test. This procedure leads to rejecting the same 3 tests that were rejected above.

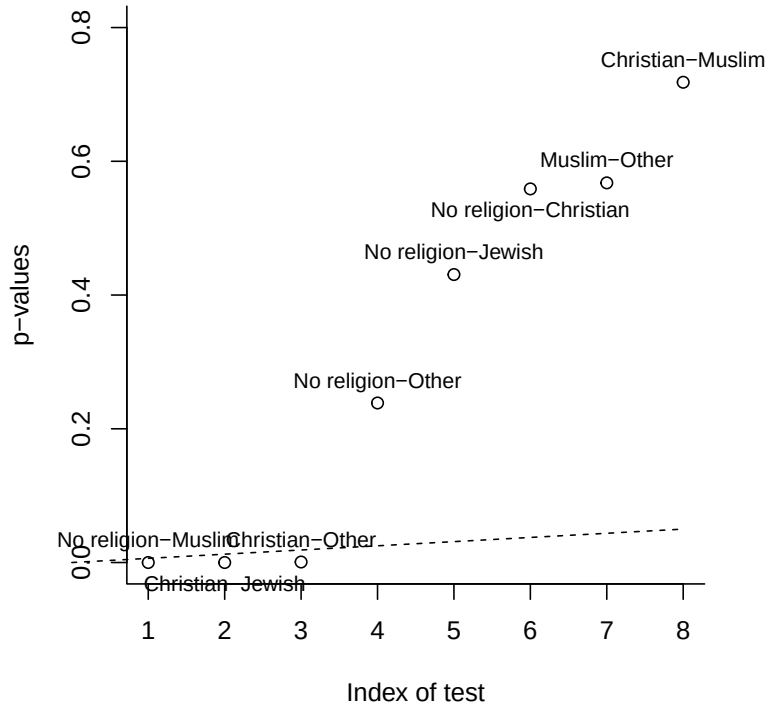

Figure 1.2: Benjamini–Hochberg procedure for multiple testing. p-values under the dashed line imply rejection of the corresponding test.

Observed vs. simulated transmission rates by religion and education categories

See Figure 1.3.

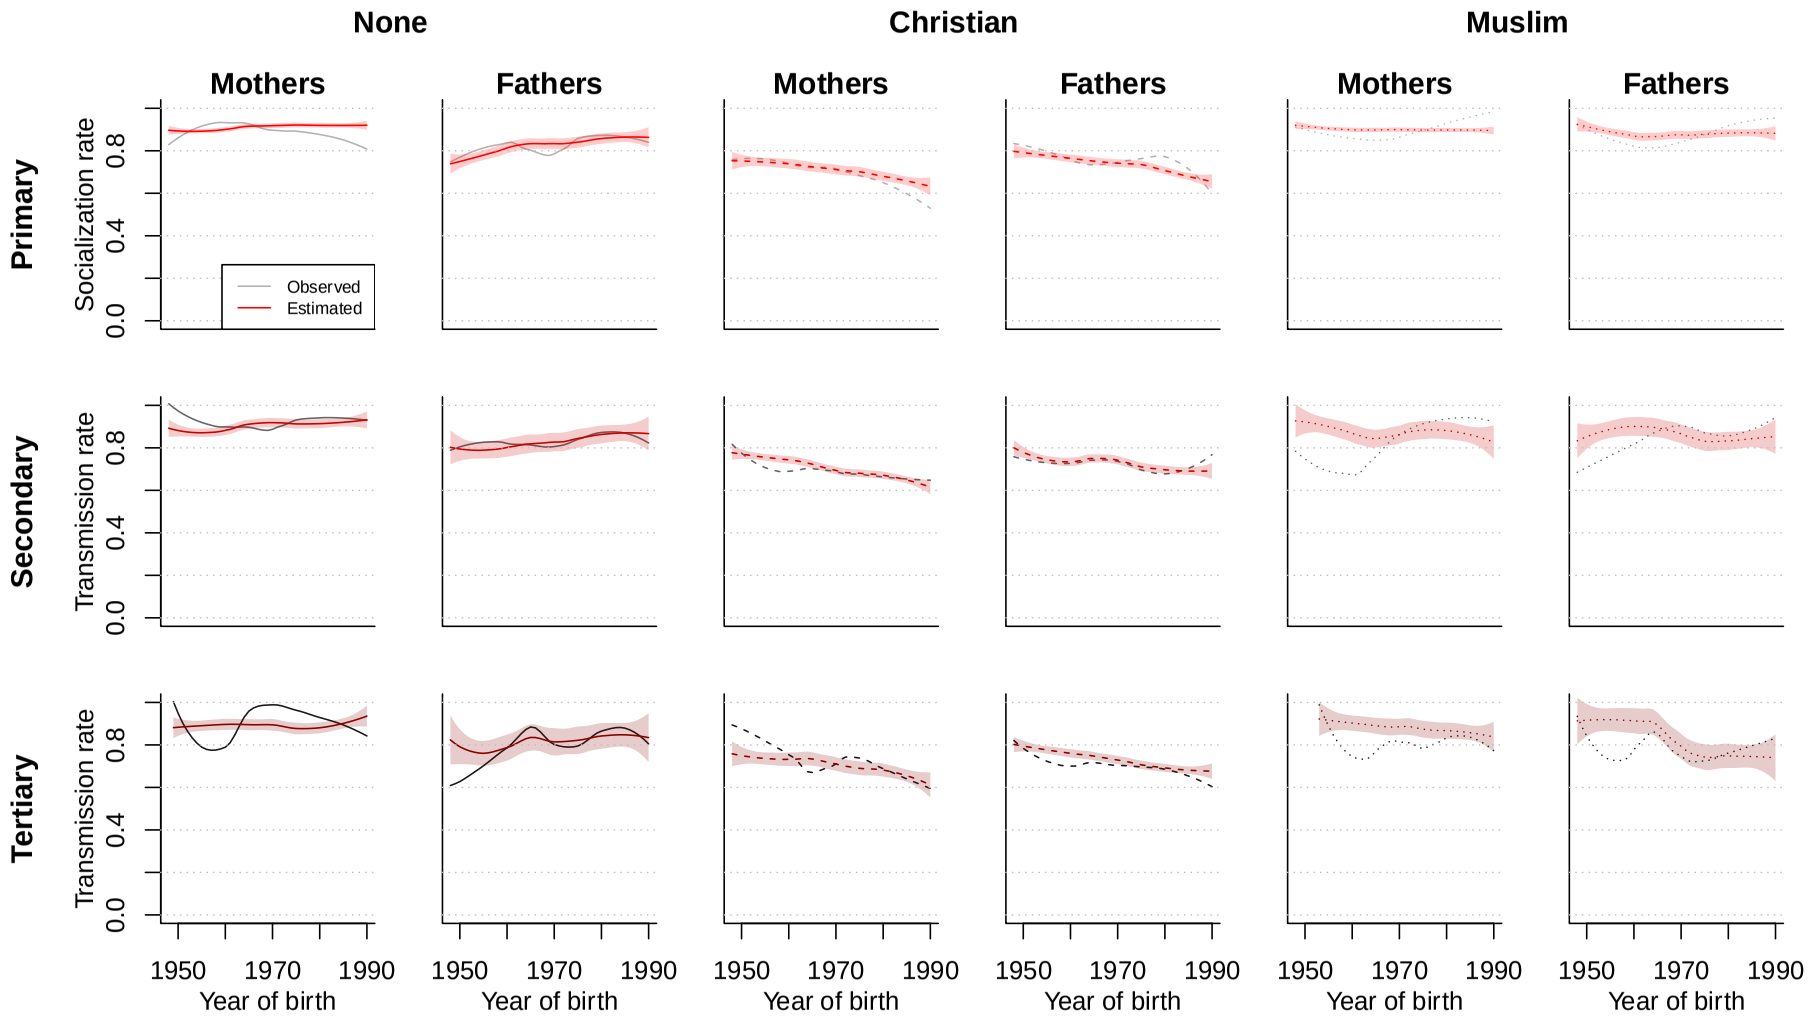

Figure 1.3: Conditional logit transmission, smoothed observed vs. simulated transmission rates (by Religion & Education).

1.7.3 Household formation and population dynamics

Household formation. The next step is to embed the collective household model into a matching framework, in which men and women match on the two characteristics {religion, education}. In the classical framework, women and men meet on a bilateral, frictionless marriage market. Households are formed endogenously based on the indirect utility provided by the match to each of the spouses. The associated equilibrium concept is stability: a matching is stable if and only if no two individuals would rather match together than stay in their current match.

The matching models usually fall into one of three categories: transferable utility (TU), imperfectly transferable utility (ITU), and nontransferable utility (NTU) (Chiappori 2017). Here, the homogamous household problem has the NTU property (under the assumption that the individual value of culture is homogeneous within a given culture), while the heterogamous household problem has the ITU property. The NTU case is well documented (Roth and Sotomayor 1990). Recent works provide both theoretical and empirical results for the ITU case (Galichon et al. 2019, Galichon and Salanié 2022).

The first step to analyze matching in the ITU framework is to describe the Pareto frontier of the household by expressing the utility of parent 1 as a decreasing function of the utility of parent 2,

$$u_1 = \Phi(\theta_1, \theta_2, u_2)$$

where $\theta_i = (n_i, h_i) \in \Theta$ is the bidimensional type of parent i, and Φ is decreasing in u_2. A match is then characterized by a measure ψ over Θ^2 and utility functions $u_1(\theta_1)$ and $u_2(\theta_2)$ such that

$$u_1(\theta_1) = \Phi(\theta_1, \theta_2, u_2(\theta_2)) \qquad \forall (\theta_1, \theta_2) \in \operatorname{supp} \psi.$$

Stability requires

$$u_1(\theta_1) \geq \Phi(\theta_1, \theta_2, u_2(\theta_2)) \qquad \forall (\theta_1, \theta_2) \in \Theta^2$$

which implies

$$u_1(\theta_1) = \max_{\vartheta_2} \Phi(\theta_1, \vartheta_2, u_2(\vartheta_2))$$

and similarly for $u_2(\theta_2)$. One can then use first-order conditions to analyze the matching problem.

In my case, even though the function Φ exists, I cannot find a closed-form expression for it.

Instead, I can parametrize the Pareto frontier by the power μ,

$$u_1 = \Phi_1(\theta_1, \theta_2, \mu)$$
$$u_2 = \Phi_2(\theta_1, \theta_2, \mu),$$

where Φ_1 is increasing and Φ_2 is decreasing in μ. A match must then be characterized by a measure ψ over Θ^2, utility functions $u_1(\theta_1)$ and $u_2(\theta_2)$, and a power function $\mu(\theta_1, \theta_2)$ such that

$$u_1(\theta_1) = \Phi_1(\theta_1, \theta_2, \mu(\theta_1, \theta_2))$$
$$u_2(\theta_2) = \Phi_2(\theta_1, \theta_2, \mu(\theta_1, \theta_2))$$

for all $(\theta_1, \theta_2) \in \operatorname{supp} \psi$. Stability requires

$$u_1(\theta_1) \geq \Phi_1(\theta_1, \theta_2, \mu(\theta_1, \theta_2))$$
$$u_2(\theta_2) \geq \Phi_2(\theta_1, \theta_2, \mu(\theta_1, \theta_2))$$

for all $(\theta_1, \theta_2) \in \Theta^2$, implying

$$u_1(\theta_1) = \max_{\vartheta_2} \ \Phi_1(\theta_1, \vartheta_2, \mu(\theta_1, \vartheta_2))$$
$$u_2(\theta_2) = \max_{\vartheta_1} \ \Phi_2(\vartheta_1, \theta_2, \mu(\vartheta_1, \theta_2)).$$

First-order conditions with respect to h_2 and h_1 write

$$\frac{\partial \Phi_1}{\partial h_2}(\theta_1, \theta_2, \mu(\theta_1, \theta_2)) + \frac{\partial \mu}{\partial h_2}(\theta_1, \theta_2) \times \frac{\partial \Phi_1}{\partial \mu}(\theta_1, \theta_2, \mu(\theta_1, \theta_2)) = 0$$
$$\frac{\partial \Phi_2}{\partial h_1}(\theta_1, \theta_2, \mu(\theta_1, \theta_2)) + \frac{\partial \mu}{\partial h_1}(\theta_1, \theta_2) \times \frac{\partial \Phi_2}{\partial \mu}(\theta_1, \theta_2, \mu(\theta_1, \theta_2)) = 0$$

which is a partial differential equation for μ.

The following issues arise compared to the usual framework:

- There is no explicit form for the function Φ that allows to describe the Pareto frontier with u_1 as a function of u_2. Consequently, I must rely on parametrizing the Pareto frontier by the Pareto weight μ, thus introducing a new function into the equilibrium. A consequence is that μ must be recovered through a system of partial differential equations rather than a standard differential equation for recovering utilities.
- The type of individuals is bidimensional, with the first dimension being discrete. Thus, the solution cannot be characterized entirely by first-order conditions.

The empirical analysis might, however, be easier. The bidimensional type is now (n, e), which takes a finite number $N \times E$ of values. Index these types by I for women and J for men. The goal is to find the Pareto weights μ^{IJ} that best explain the empirical matching patterns, according to the individuals' discrete choices. Denoting women by $i \in I$ and men by $j \in J$, these discrete choice problems are

$$u_i = \max_j \left\{ \Phi_1(I, J, \mu^{IJ}) + \alpha_i^J \right\}$$
$$u_j = \max_i \left\{ \Phi_2(I, J, \mu^{IJ}) + \beta_j^I \right\}$$

where α_i^J and β_j^I are random shocks that depend exclusively on the partner's type, as in the Choo and Siow (2006) framework. These translate into a probability for each individual i or j of marrying a partner of type J or I. In this case, estimation must be performed simultaneously on marriage patterns and transmission patterns to jointly estimate the Pareto weights μ^{IJ} and the parameters of the utilities and production functions that I estimated previously.

Population dynamics. The last contribution of this paper will be to study the population dynamics implied by the model. This can be conducted either empirically or theoretically. Empirically: once the model's primitive parameters are estimated, one can iterate the model to simulate the evolution of the population along the two dimensions of interest (religion and education). This simply requires to solve for the ITU matching equilibrium, for which a solution was proposed by Galichon and Salanié (2022). From the matching equilibrium, we can infer the joint distribution of religious traits and educational levels in the next generation through the collective household model. It might be possible to perform the same exercise theoretically for sufficiently simple distributions of traits.

The dynamic implications could be interesting. For instance, if the cultural minority starts with lower average human capital than the majority (as is for instance the case with immigrants in many countries), the need to safeguard their culture could occur at the expense of their human capital development, such that the human capital gap between cultural minority and majority could widen with time. (Or, at the least, this mechanism could delay the catch-up of the minority with the majority compared to the baseline case wherein people do not care about cultural transmission.) Intuitively, this process could lead to a higher-educated, little-socialized cultural majority on the one side and a lower-educated, highly socialized cultural minority on the other side.

1.7.4 Analysis of deviance residuals

In order to examine the validity of the hypothesis of independent errors in section 1.5.1, here I analyze the residuals of the estimated structural model. In qualitative response models such as multinomial logit or ordered logit (which are the two models that I use), there are several options for computing residuals. Notable examples include response residuals, Pearson residuals, generalized residuals, or deviance residuals. Deviance residuals, in particular, are obtained by measuring the contribution of each individual observation to the total deviance of the estimated model. In models with multiple choice they are the easiest to handle because they are one-dimensional – whereas in the case of Pearson or generalized residuals, there are as many residuals as there are possible responses. For this reason, i choose deviance residuals for this analysis.

From equation (1.26), we can rewrite the deviance of the model as

$$-2 \ln L = \sum_i w_i \left(d_i^{\text{rel}} + d_i^{\text{edu}} \right) \tag{1.35}$$

where

$$d_i^{\text{rel}} = -2 \sum_{n=1}^{N} \mathbf{1}_{\{i \text{ is } n\}} \ln(\pi_{in}) \qquad \text{and} \qquad d_i^{\text{edu}} = -2 \sum_{e=1}^{E} \mathbf{1}_{\{i \text{ is } e\}} \ln(\phi_{ie}) \tag{1.36}$$

are the contributions of the individual observation i to the deviance, in terms of religious affiliation (d_i^{rel}) and educational attainment (d_i^{edu}) respectively. Deviance residuals r_i^{rel} and r_i^{edu} are then defined as

$$r_i^{\text{rel}} = (-1)^{\mathbf{1}\{n_i \neq \arg\max_n K_{in}\}} \sqrt{d_i^{\text{rel}}} \tag{1.37}$$

$$r_i^{\text{edu}} = (-1)^{\mathbf{1}\{\ln H_i > \bar{h}_{e_i}\}} \sqrt{d_i^{\text{edu}}}. \tag{1.38}$$

To understand the signs, recall that n_i is the observed religion of i and K_{in} her predicted level of religious capital in religion n. The condition $n_i \neq \arg\max_n K_{in}$ is then satisfied when the religious affiliation predicted by the model for i is different than the actual one. Thus, r_i^{rel} is positive if the model correctly predicted the religious affiliation of individual i, and negative otherwise. This sign is consistent with the definition of response residuals for binomial logit models, for instance.

For the education residuals, recall that e_i is the observed education level of i, and $\bar{h}_{e_i}$ is the ordered logit threshold between having education level e_i and $e_i + 1$. Furthermore, $\ln H_i$ is the predicted level of (log-)human capital for i. Thus, r_i^{edu} is positive if the model predicted an education level identical or below the observed one, and negative otherwise. This is consistent with the common understanding of residuals (e.g. in traditional linear models), in which residuals

are negative if the model "overshoots," and positive if it "undershoots."

I compute these deviance residuals, and present them in Figure 1.1. The plot represents the education residuals r_i^{edu} as a function of the religion residuals r_i^{rel}, as well as the best linear prediction. Note that there are no residuals r_i^{rel} between -1 and 0 (roughly). This is a mechanical consequence of the multinomial logit model: if an affiliation n is not predicted by the model (negative residuals), it means that its associated choice probability must be below $1/2$ (otherwise it would be the most likely outcome, i.e. the predicted outcome). As a consequence, in this case the deviance contribution d_i^{rel} must be more than $-2\ln(1/2) = 2\ln 2$. Finally, the residual must be less than $-\sqrt{2\ln 2} \simeq 1.18$; this is consistent with the observed values for r_i^{rel}. For a similar reason, there are no residuals r_i^{edu} between -1 and 0 (again, roughly), leading to an "empty cross" pattern.

The linear fit suggests a very weak negative correlation between the religious affiliation residuals r_i^{rel} and the educational attainment residuals r_i^{edu}. This suggests that the assumption on the independence of errors is reasonable.

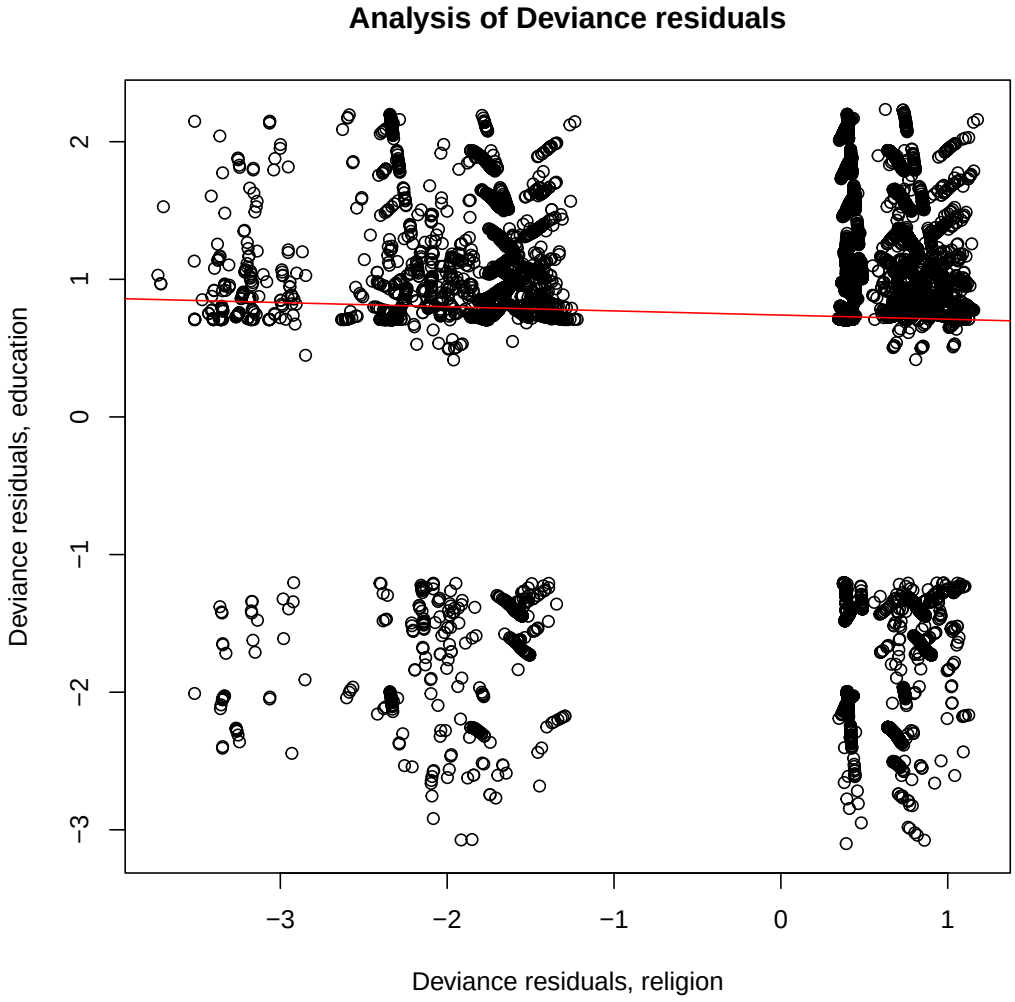

Figure 1.1: Deviance residuals on religion r_i^{red} and education r_i^{edu}, along with linear prediction (in red).

1.7.5 Educational homogamy: local log odds ratios analysis

In this section I follow the methodology of Siow (2015) to study educational homogamy using local log odds ratios. As pointed out by Siow, simply computing correlations of spouses' education levels remains a weak test of homogamy since we don't know how high the correlation should be to infer that the data indeed exhibits homogamy. A stronger tests consists in verifying that all local log odds ratios are positive.

To begin with, Table 1.1 provides the sample distribution of marriages according to the spouses' education levels. The repartition of marriages is thus represented by a 3×3 matrix $(n_{ij})_{1 \leq i,j \leq 3}$, for a total number of observations N. The local log odds ratios are defined for $i, j \leq 2$ as

$$\ln \left(\frac{n_{ij}\, n_{i+1,j+1}}{n_{i,j+1}\, n_{i+1,j}} \right) \tag{1.39}$$

which constitutes a measure of local homogamy in the submatrix $\left(\begin{smallmatrix} n_{ij} & n_{i,j+1} \\ n_{i+1,j} & n_{i+1,j+1} \end{smallmatrix}\right)$. In particular if random matching is occurring, one should expect all these log odds ratios to be equal to 0.

Siow (2015) shows that supermodularity of the marital surplus implies that all local log odds ratios should be positive, i.e. that the matrix $(n_{ij})_{1\leq i,j\leq 3}$ should be totally positive of order 2, or TP2 for short. I test this TP2 criterion statistically by following the method prescribed by Garre et al. (2002), which Siow (2015) also follows. First define three different hypotheses: H_0 corresponds to the restricted model where all local log odds ratios are equal to 0; H_1 the model where they are positive; and H_2 the unrestricted model. Hypothesis H_0 also means that the matrix $(n_{ij})_{1\leq i,j\leq 3}$ is totally null of order 2, which I call TN2 for short. Call L_0, L_1, and L_2 the models' respective likelihoods: for instance,

$$L_1 = \max_{\nu_{ij}} \sum_{ij} n_{ij} \ln(\nu_{ij}) \tag{1.40}$$

subject to the constraints

$$\ln\left(\frac{\nu_{ij}\ \nu_{i+1,j+1}}{\nu_{i,j+1}\ \nu_{i+1,j}}\right) \geq 0 \qquad (\forall i,j \leq 2) \tag{1.41}$$

and

$$\sum_{ij} \nu_{ij} = N. \tag{1.42}$$

The likelihood L_0 is obtained by using an equality constraint in (1.41), and L_2 by removing constraint (1.41) entirely. The statistics of interest are log-likelihood ratio (LR) test statistics,

$$\mathrm{LR}_{01} = 2(L_1 - L_0) \qquad \text{and} \qquad \mathrm{LR}_{12} = 2(L_2 - L_1). \tag{1.43}$$

The statistic LR_{12} indicates to what extent TP2 fits the data, and LR_{01} tests whether positive local log odds ratios are a better fit than if they are null. When samples obey TP2, I test H_1 versus H_0. When they do not, I test H_1 versus H_2. I report estimates of the probabilities

Mother's education	Father's education			Total
	Primary or less	Secondary	More than secondary	
Primary or less	8998	1968	298	11264
Secondary	1136	3428	1023	5587
More than secondary	97	428	1417	1942
Total	10231	5824	2738	18793

Table 1.1: Parental education and homogamy.

	Father's education					Local log odds ratios	
Mother's education	Pri	Sec	Sec+	Total		Pri, Sec	Sec, Sec+
	TP2 probabilities					TP2 log odds	
Pri	0.479	0.105	0.016	0.600	Pri, Sec	2.624	0.678
	(0.004)	(0.002)	(0.001)			(0.039)	(0.068)
Sec	0.060	0.182	0.054	0.296	Sec, Sec+	0.380	2.406
	(0.002)	(0.003)	(0.002)			(0.119)	(0.061)
Sec+	0.005	0.023	0.075	0.103			
	(0.001)	(0.001)	(0.002)		LR_{01} statistic: 10 352		
Total	0.544	0.310	0.145	1	p-value: 0		
	TN2 probabilities						
Pri	0.326	0.186	0.087	0.599			
	(0.003)	(0.002)	(0.002)				
Sec	0.162	0.092	0.043	0.297			
	(0.002)	(0.001)	(0.001)				
Sec+	0.056	0.032	0.015	0.103			
	(0.001)	(0.001)	(0.000)				
Total	0.544	0.310	0.145	1			

Note: Standard errors in parentheses (parametric bootstrap, 100 replications).

Table 1.2: Estimated probabilities and local log odds ratios – full sample.

$p_{ij} = \frac{\nu_{ij}}{N}$ for a marriage observation to fall in the ij category. The p-values and standard errors are obtained by parametric bootstrap with 100 replications.

Analysis on the full sample. Table 1.2 presents the estimated probabilities and the associated local log odds ratios for the full sample. The local log odds ratios are all positive, so the data obeys TP2. For this reason, the estimates from the unrestricted problem are the same as the TP2 estimates, which is why I only report the latter. In this case, the relevant hypothesis test is H_1 versus H_0: is there evidence for positive local log odds ratios, rather than them being all zeros? The associated test statistic is LR_{01}.

The value of the LR_{01} test statistic is very large in this case, at 10 352. Accordingly, the p-value is extremely small – in fact, it cannot be differentiated from 0 at the precision level which I use. This provides strong evidence to reject the null hypothesis H_0 that local log odds ratios are all zeros, in favor of H_1 and TP2. In turn, this provides strong evidence of homogamy and of the supermodularity of the marital surplus in the full sample.

Analysis conditional on spouses' religious affiliations. Figure 1.1 presents the empirical local log odds ratios conditional on spouses' religious affiliation using a color chart. Red indicates positive values, and blue negative ones (gray indicates missing data). A glimpse at the figure

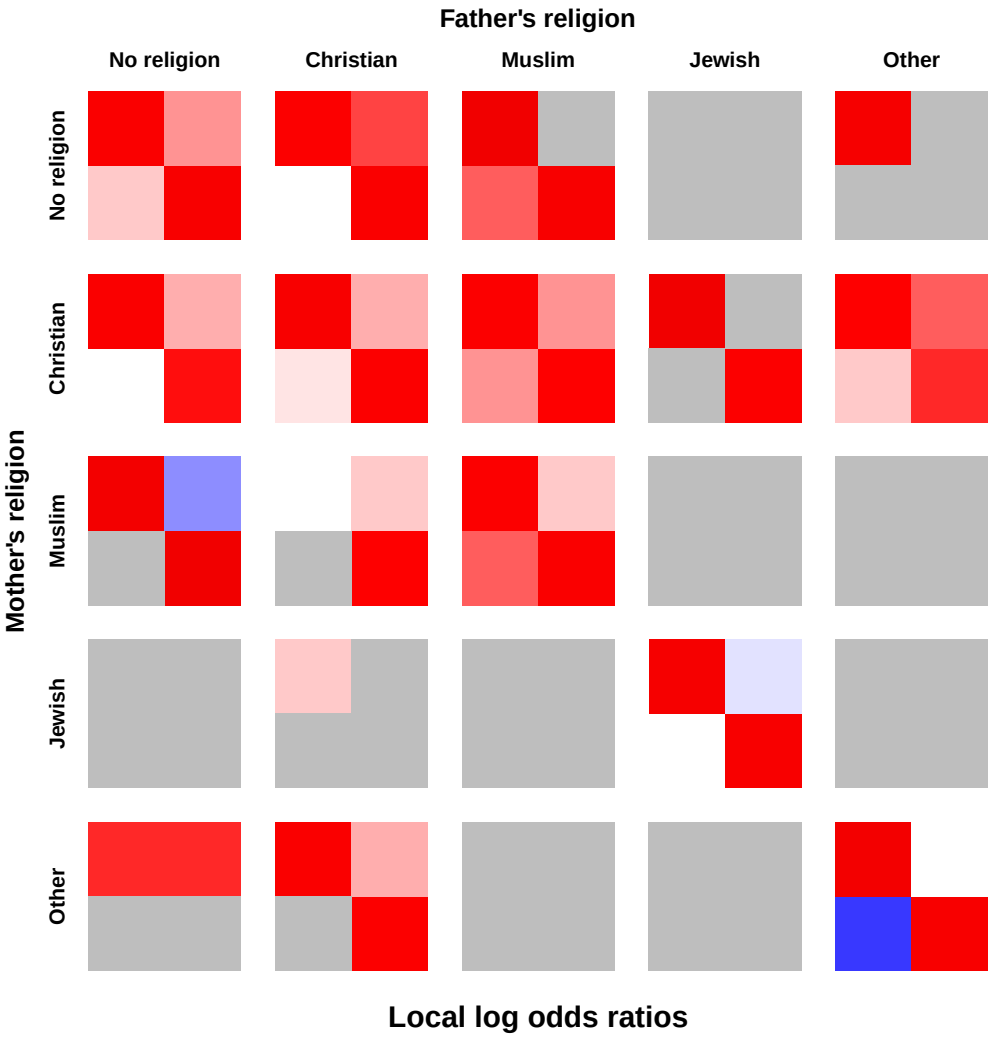

Figure 1.1: Local log odds ratios conditional on spouses' religious affiliation. Red indicates positive values, and blue negative ones. Lighter shades indicate values closer to 0. Gray indicates missing values.

shows that most of the local log odds ratios which can be computed are positive. I test TP2 for each configuration of the spouses' religious affiliation, using the same method as for the full sample.

Mother: No religion				Father: No religion		
	Father's education				Local log odds ratios	
Mother's education	Pri	Sec	Sec+		Pri, Sec	Sec, Sec+
	Unrestricted probabilities				Unrestricted log odds	
Pri	0.342	0.094	0.009	Pri, Sec	2.425	0.860
	(0.011)	(0.007)	(0.002)		(0.134)	(0.300)
Sec	0.085	0.264	0.059	Sec, Sec+	0.531	2.657
	(0.006)	(0.010)	(0.006)		(0.322)	(0.162)
Sec+	0.006	0.034	0.107			
	(0.002)	(0.004)	(0.007)			
	TP2 probabilities				TP2 log odds	
Pri	0.342	0.094	0.009	Pri, Sec	2.425	0.860
	(0.011)	(0.007)	(0.002)		(0.119)	(0.301)
Sec	0.085	0.264	0.059	Sec, Sec+	0.531	2.657
	(0.006)	(0.009)	(0.005)		(0.318)	(0.161)
Sec+	0.006	0.034	0.107			
	(0.002)	(0.004)	(0.006)			
	TN2 probabilities			$N = 2033$		
Pri	0.193	0.175	0.078	LR_{12} statistic: 0		
	(0.007)	(0.005)	(0.004)	p-value: 1		
Sec	0.176	0.160	0.071	LR_{01} statistic: 1 211		
	(0.005)	(0.005)	(0.004)	p-value: 0		
Sec+	0.064	0.058	0.026			
	(0.004)	(0.004)	(0.002)			

Note: Standard errors in parentheses (parametric bootstrap, 100 replications).

Table 1.3: Estimated probabilities and local log odds ratios – No religion, No religion.

Mother: No religion				Father: Christian		
	Father's education				Local log odds ratios	
Mother's education	Pri	Sec	Sec+		Pri, Sec	Sec, Sec+
	Unrestricted probabilities				Unrestricted log odds	
Pri	0.175	0.079	0.005	Pri, Sec	2.212	1.427
	(0.026)	(0.021)	(0.005)		(0.462)	(38.636)
Sec	0.069	0.286	0.079	Sec, Sec+	0.042	2.454
	(0.020)	(0.028)	(0.017)		(7.964)	(0.404)
Sec+	0.016	0.069	0.222			
	(0.010)	(0.018)	(0.027)			
	TP2 probabilities				TP2 log odds	
Pri	0.175	0.079	0.005	Pri, Sec	2.212	1.427
	(0.031)	(0.021)	(0.005)		(0.496)	(23.882)
Sec	0.069	0.286	0.079	Sec, Sec+	0.042	2.454
	(0.019)	(0.033)	(0.017)		(13.175)	(0.389)
Sec+	0.016	0.069	0.222			
	(0.007)	(0.019)	(0.030)			
	TN2 probabilities			$N = 189$		
Pri	0.067	0.112	0.080	LR_{12} statistic: 0		
	(0.011)	(0.015)	(0.013)	p-value: 1		
Sec	0.112	0.188	0.133	LR_{01} statistic: 109		
	(0.015)	(0.022)	(0.017)	p-value: 0		
Sec+	0.080	0.133	0.094			
	(0.012)	(0.019)	(0.015)			

Note: Standard errors in parentheses (parametric bootstrap, 100 replications).

Table 1.4: Estimated probabilities and local log odds ratios – No religion, Christian.

Mother: No religion	Father: Muslim					
	Father's education				Local log odds ratios	
Mother's education	Pri	Sec	Sec+		Pri, Sec	Sec, Sec+
	Unrestricted probabilities				Unrestricted log odds	
Pri	0.375	0.011	0.000	Pri, Sec	3.561	+Inf
	(0.057)	(0.012)	(0.000)		(20.911)	(20.160)
Sec	0.170	0.182	0.045	Sec, Sec+	1.322	2.639
	(0.038)	(0.042)	(0.022)		(22.206)	(9.429)
Sec+	0.011	0.045	0.159			
	(0.012)	(0.024)	(0.043)			
	TP2 probabilities				TP2 log odds	
Pri	0.375	0.011	0.000	Pri, Sec	3.561	26.579
	(0.049)	(0.013)	(0.000)		(13.546)	(5.671)
Sec	0.170	0.182	0.045	Sec, Sec+	1.322	2.639
	(0.043)	(0.045)	(0.021)		(14.614)	(0.914)
Sec+	0.011	0.045	0.159			
	(0.011)	(0.022)	(0.043)			
	TN2 probabilities			$N = 88$		
Pri	0.215	0.092	0.079	LR_{12} statistic: 0		
	(0.036)	(0.022)	(0.022)	p-value: 1		
Sec	0.221	0.095	0.081	LR_{01} statistic: 71		
	(0.042)	(0.023)	(0.019)	p-value: 0		
Sec+	0.120	0.052	0.044			
	(0.027)	(0.017)	(0.013)			

Note: Standard errors in parentheses (parametric bootstrap, 100 replications).

Table 1.5: Estimated probabilities and local log odds ratios – No religion, Muslim.

Mother: No religion	Father: Jewish					
	Father's education				Local log odds ratios	
Mother's education	Pri	Sec	Sec+		Pri, Sec	Sec, Sec+
	Unrestricted probabilities				Unrestricted log odds	
Pri	0.250	0.000	0.000	Pri, Sec	–	–
	(0.163)	(0.000)	(0.000)		(9.925)	(11.615)
Sec	0.375	0.000	0.125	Sec, Sec+	–	–
	(0.190)	(0.000)	(0.115)		(4.953)	(14.785)
Sec+	0.000	0.000	0.250			
	(0.000)	(0.000)	(0.162)			
	TP2 probabilities				TP2 log odds	
Pri	0.250	0.000	0.000	Pri, Sec	7.410	13.921
	(0.132)	(0.000)	(0.000)		(3.095)	(5.022)
Sec	0.375	0.000	0.125	Sec, Sec+	15.385	8.021
	(0.161)	(0.000)	(0.123)		(4.501)	(8.117)
Sec+	0.000	0.000	0.250			
	(0.000)	(0.000)	(0.144)			
	TN2 probabilities			$N = 8$		
Pri	0.156	0.000	0.094	LR_{12} statistic: 0		
	(0.105)	(0.000)	(0.079)	p-value: 0.750		
Sec	0.312	0.000	0.187	LR_{01} statistic: 6.086		
	(0.139)	(0.000)	(0.103)	p-value: 0.010		
Sec+	0.156	0.000	0.094			
	(0.114)	(0.000)	(0.074)			

Note: Standard errors in parentheses (parametric bootstrap, 100 replications).

Table 1.6: Estimated probabilities and local log odds ratios – No religion, Jewish.

Mother: No religion	Father: Other					
	Father's education				Local log odds ratios	
Mother's education	Pri	Sec	Sec+		Pri, Sec	Sec, Sec+
	Unrestricted probabilities				Unrestricted log odds	
Pri	0.250	0.000	0.000	Pri, Sec	–	–
	(0.163)	(0.000)	(0.000)		(9.925)	(11.615)
Sec	0.375	0.000	0.125	Sec, Sec+	–	–
	(0.190)	(0.000)	(0.115)		(4.953)	(14.785)
Sec+	0.000	0.000	0.250			
	(0.000)	(0.000)	(0.162)			
	TP2 probabilities				TP2 log odds	
Pri	0.250	0.000	0.000	Pri, Sec	7.410	13.921
	(0.132)	(0.000)	(0.000)		(3.095)	(5.022)
Sec	0.375	0.000	0.125	Sec, Sec+	15.385	8.021
	(0.161)	(0.000)	(0.123)		(4.501)	(8.117)
Sec+	0.000	0.000	0.250			
	(0.000)	(0.000)	(0.144)			
	TN2 probabilities			$N = 24$		
Pri	0.156	0.000	0.094	LR_{12} statistic: 0		
	(0.105)	(0.000)	(0.079)	p-value: 1		
Sec	0.312	0.000	0.187	LR_{01} statistic: 25.125		
	(0.139)	(0.000)	(0.103)	p-value: 0		
Sec+	0.156	0.000	0.094			
	(0.114)	(0.000)	(0.074)			

Note: Standard errors in parentheses (parametric bootstrap, 100 replications).

Table 1.7: Estimated probabilities and local log odds ratios – No religion, Other.

Chapter 2

Veiling and the Economic Integration of Muslim Women in France

This chapter was co-written with Sébastien Montpetit (Toulouse School of Economics).

2.1 Introduction

Veiling among Muslim women has been at the center of public debates in Western countries for several decades. The Islamic veil is often perceived as a signal of both cultural distance from the majority, and of the subordination of women. It is a particularly burning issue in France, where state secularism (*laïcité*) "constitutes a pillar, even the identity and foundation of the community life."[1] At the heart of the debates lies the idea that Muslim women wear the veil against their own will and must be freed from such oppression.

To be sure, the adoption of this cultural practice entails numerous costs such as reduced employment prospects, discrimination, and physical discomfort (Abdelhadi 2019, Valfort 2020). Yet, as many politicians advocate for a strengthening of secular policies, it is crucial to understand the real motives behind veiling, and how it affects the economic participation of Muslim women. First, do women veil willingly despite these costs, or is veiling mostly a result of communitarian pressures? The answer to this question may lead to opposite policy recommendations: if veiling is driven by individual motives, then further restrictions on veiling may inhibit the socio-economic integration of Muslim women even more and reduce social welfare (Carvalho 2013, Shofia 2020). But if veiling is community-driven, then those restrictions may help emancipate them (Maurin and Navarrete-Hernandez 2023). Second, are veiled women less economically active because of

[1] Andriantsimbazovina et al. (2020), p. 7.

religious preferences, or because they face more obstacles in the labor market? If the latter is true, the objective of policymakers who wish to improve the economic participation of veiled women shouldn't be to ban the veil in the workplace, thus alienating Muslim women from the workforce even further, but instead to find ways to remove the barriers to economic integration that veiled women are already facing.

Despite the considerable media, political, and academic attention, the reasons why women veil in a secular and Muslim-minority country like France are still poorly understood. This is in contrast with the context of Muslim-majority countries, which has received attention both in economics (Carvalho 2013, Shofia 2020) and in the wider social science literature. In Muslim-minority countries, most of the empirical evidence on veiling behavior remains based on interviews conducted over small samples of women (or adolescents). Moreover, in France, such interviews are typically conducted in the Parisian region, even though Muslims are increasingly present over the whole territory. In addition to this representativeness issue, this methodology has the inherent drawback that, especially for such sensitive topics, interviewees may be susceptible to social desirability bias. This is even more true because respondents are typically aware that the topic of the interview is veiling behavior. It is thus not clear how individuals' responses reflect true individual preferences for veiling or influences from their community.

In this paper, we make one of the first attempts at analyzing the relationship between veiling and economic participation, using rich survey data over more than 3,000 Muslim women in France. This sample constitutes the largest source of data on Muslim women in France and their veiling practices that we are aware of. In addition, its wide geographical coverage arguably improves representativeness compared to interview-based data. The survey also records other detailed information about respondents, providing important controls for our analysis. Overall, these data allow us to study veiling and economic participation among Muslim women empirically on a scale which hasn't been done in a Muslim-minority country before. Furthermore, this paper also extends the existing economic theory of veiling to the context of Muslim-majority countries. The structure of the model notably helps us to disentangle the role of religious motives versus those of economic motives in women's veiling and economic participation decisions.

A second objective of this paper is to unpack the various motives for veiling. By matching our main data with other sources, we are able to measure the influence of the local community on women's veiling and economic participation. We also exploit the richness of the survey to proxy for parental religious transmission, individual religiosity, and the individual's religious environment.

Our study begins with an in-depth descriptive analysis, where we provide evidence that in

France, wearing conspicuous religious symbols is associated with much lower levels of economic participation. Using a rich set of controls, we notably find that the practice of always wearing such a symbol in public is associated with a decline of 23 percentage points in economic participation (defined as being active on the labor market or studying) in the cross-section. This correlation is large and economically significant. In our preferred specification for instance, veiling is associated with a decline in economic participation which is equivalent to having an additional 1.4 children aged less than 4 years old. We find that this negative relationship is robust to several alternate specifications. In particular, exploiting the information on respondents' employment history, we construct a retrospective panel dataset of economic participation. We show that the estimated negative correlation is robust to the inclusion of year fixed effects and random effects and is similar in magnitude to that obtained in the cross-section.

In a second step, we develop a model to analyze the joint decision of veiling and economic participation. Our goal is to provide a conceptual framework to understand the respective roles of religious motives (such as individual religiosity or religious social pressure) and of economic motives (such as employment opportunities and on-the-job discrimination) in this joint decision. The model nests Carvalho's (2013) seminal theory of veiling, but also extends it to fit the French context based on our descriptive results and on our understanding of the ethnographic evidence. In the original theory, veiling is a response to individual and social religious incentives: veiling acts as a commitment device to follow religious norms and as a signal of the woman's commitment to her community. In addition to this religious incentives channel, we introduce economic incentives to (un)veil in the model, which reflect the documented barriers to economic participation that veiled women face. These two mechanisms have different implications for how the decisions to veil and to participate economically interact: according to the religious incentives channel, women should veil more when they participate (in order to signal their religious commitment despite their social integration), while according to the economic incentives channel, they should veil less (because veiling directly reduces their economic opportunities).

Finally, we translate our conceptual framework into an empirical static discrete-choice model of veiling and economic participation. We formulate and test direct implications of the theory for the religious incentives and economic incentives channels. For the religious incentives channel, we distinguish between intrinsic motives and social religious pressure in the joint decision. We measure the intrinsic motives using multiple indicators of religiosity (of both subjective feelings and actual religious practices) available in the survey data. For social religious pressure, we develop several proxies. Parental influence is measured using the (self-reported) importance of religion in the education received by the respondent and religious name-giving. For communi-

tarian pressure, given that data on religious diversity is not available in France, we use the share of Maghrebi immigrants in the local population as well as the local number and size of Muslim places of worship (mosques and prayer rooms).

Our main empirical findings are twofold. First, we find supporting evidence for the economic discrimination channel described in the theory, but not for the religious incentives channel. This result suggests that the impact of religious motives on the economic participation decision is mostly indirect (through the decision to veil), while economic motives seem to have a direct impact on the decision to veil. In other words, the primary reason why veiled Muslim women work less (or, equivalently, that working Muslim women veil less) seems to be that veiling itself reduces their economic opportunities, and not that religiosity disincentivizes working. As such, the lower economic participation of Muslim women could be understood as a demand-side problem on the labor market, more than a supply-side one.

Second, we measure the respective roles of the different religious motives in the decision to veil. While measures of social religious pressures are correlated with veiling behavior, we find that a much larger share of the variation in veiling patterns can be explained by individual religiosity. Our results thus question the rhetoric often used to justify policies restricting the wearing of religious symbols in France. Consistent with our analytical results, we conjecture that regulations which limit the expression of religious faith in public are likely to impede integration of Muslim women into Western societies.

2.1.1 Related literature and contributions

This paper contributes to several strands of the literature. First, it provides novel empirical evidence to the vast literature on Islamic veiling in the social sciences.[2] In this literature, most of the evidence is based on interviews with Muslim women since veiling behavior is rarely observed in surveys or other standard datasets. While interviews have the potential to dig deeper into specific questions of interest and uncover a large number of potential channels, they often suffer from small sampling and representativeness issues. In a recent contribution, Shofia (2020) measures the veiling rate at the district level to circumvent this problem and provides robust empirical evidence that better economic opportunities for women induce Indonesian women to veil. In contrast, in this paper, we study the case of a secular country in which Muslims form a minority and where wearing the veil is frowned upon rather than encouraged. Similar conclusions to that of Shofia (2020) were reached by Aksoy and Gambetta (2016), the closest study to ours,

[2]We review in detail the literature on veiling in France in section 2.2. Recent contributions in other contexts include Harrison (2016) for the United States as well as Aksoy (2017) and Aksoy and Gambetta (2016, 2021) for Turkey.

for the case of Turkey. Aksoy and Gambetta (2016) also attempt to study the determinants of veiling in a Western country, namely Belgium. However, they do not have a direct measure of veiling behavior , but rather a measure of attitudes towards veiling in public. Moreover, the richness of our data allows us to further unpack the relative weight of various incentives that are difficult to measure in the decision to wear the Islamic veil over a large sample. In particular, we can distinguish between private and communitarian incentives to veil, a question which has so far eluded empirical researchers. Another close study is that of Abdelhadi (2019) who finds that the wearing of the veil is associated with lower employment in the United States, but does not investigate the motives for veiling. Her result is consistent with our findings for France for which we document large differences in economic participation between veiled and non-veiled women.

Second, we bring new evidence on motives for adopting costly cultural practices both theoretically and empirically. In the vast literature in on the economics of religion and identity, it is now acknowledged that individuals may choose their identity via rational decision-making even if it requires costly investments or sacrifices (Iannaccone 1992, Akerlof and Kranton 2000, Atkin et al. 2021, Jia and Persson 2021). However, though potentially rational, adopting (or transmitting) certain cultural practices can be an impediment to social and economic integration of certain groups. A strand of the literature has investigated the incentives that might justify such choice. Recent examples include foot-binding in China (Fan and Wu 2022), female genital cutting in Africa (Bellemare et al. 2015, Novak 2020, Gulesci et al. 2021), and baby-naming choice in France (Algan et al. 2022).[3] We contribute to this literature in three ways. First, we document that in France, veiling is associated with poorer economic integration of Muslim women rather than being an integration strategy as suggested by evidence in Muslim-majority countries (Aksoy and Gambetta 2016, Shofia 2020). Second, we provide detailed descriptive evidence of why Muslim women might wear such a costly signal of religious identity in France. Third, we uncover novel empirical patterns concerning the wearing of *discreet* signs of religious affiliation, which have received little attention in the literature. In particular, they appear to be worn by Muslim women who are educated and moderately religious. These patterns might suggest discreet symbols, in the French context, play a similar role to that of the veil in Muslim-majority countries.[4]

Third, our results have implications for State secularisation policies. Of particular interest in

[3]There is also a relevant literature looking at incentives to abandon certain costly cultural traits and adopting less harmful ones. For example, Biavaschi et al. (2017) find important economic payoffs for the Americanization of migrants' names. See also Bisin et al. (2011, 2016) and Drydakis (2013) on economic returns of assimilation for migrants.

[4]We, however, have little statistical power to test this hypothesis because few Muslim women wear only discreet symbols in our sample.

our context, two recent empirical studies reach opposite conclusions on the effects of the French headscarf ban in public schools. On the one hand, Abdelgadir and Fouka (2020) find that the 2004 ban depressed schooling outcomes of French girls of North-African origin.[5] On the other hand, Maurin and Navarrete-Hernandez (2023) obtain that the 1994 ministerial circular asking school principals to prohibit the wearing of the veil in schools had a positive impact on their educational attainment. Even if they are comparing different cohorts of adolescents and different treatments, these contradictory pieces of evidence are puzzling. By focusing on why Muslim women are willing to sacrifice economic opportunities to veil, we can offer a new perspective to this debate. If incentives to veil are mainly *private*, more stringent secular regulations should reduce incentives to integrate for religious women who wish to veil. On the contrary, if *communitarian* incentives prevail, such veil bans may help women emancipate and liberate them from a costly religious norm which limits their economic opportunities. Our results lend support to the former interpretation. The main observed drivers of veiling behavior in France appear to be the woman's religiosity as well as non-religious identity such as her origins. Religious pressures from women's close community are also correlated with veiling behavior, but turn out to explain only a small share of variation in veiling behavior in our regressions. Proponents of French secular regulations often base their arguments on the idea that Muslim women simply do not want to veil and are forced to do so by other Muslims. Our analysis thus casts serious doubts on this assumption and suggests that the French secular regulations most likely inhibit social and economic integration of Muslim women in France rather than facilitating their emancipation.

The rest of the article is structured as follows. Section 2.2 describes the institutional context. Section 2.3 describes the data sources and provides a detailed descriptive analysis of veiling patterns in France. In section 3.2 we outline our theoretical framework. In section 2.5 we translate this framework into an empirical model and estimate its main parameters. Finally, section 3.6 concludes.

2.2 Historical and sociological background

The wearing of the Islamic veil has been a burning issue in France since at least three decades. In 1989, the "*affaire des foulards*" (headscarf affair) garnered nationwide attention when three girls were expelled from their middle school for refusing to remove their headscarves. The incident sparked heated debates but eventually culminated in the highest French administrative court ruling in favor of the expelled girls (Scott 2009). Despite this ruling, in 1994 the Ministry

[5] In a similar spirit, Benzer (2022) finds that the re-introduction of Islamic schools, which do not prohibit the headscarf, had positive impacts on girls' educational attainment in Turkey.

of Education issued a circular asking school principals to prohibit conspicuous religious symbols worn by students. This controversial position was later enshrined in a 2004 law, whose supporters argued that headscarves "infringed on the liberty of conscience of other pupils and represented the triumph of communitarian pressures" (Abdelgadir and Fouka 2020 p. 4). The debate then shifted to other public spaces, with a nationwide ban of full-face veils (*burqa*) in 2010, and later with several city bans of the *burkini* in swimming areas and beaches.[6]

Despite the significance of these policies for Muslim women and girls, they have largely been excluded from the conversation. In fact, this "one-sided debate"[7] has revealed a lack of understanding among policymakers about the realities and constraints faced by the Muslim population (Scott 2009, Nordmann 2004). Nevertheless, considerable research in sociology and anthropology has been dedicated to understanding the experience of Muslims in France, and particularly the reasons for women to wear the veil. In the following paragraphs we focus on two factors which have been shown to be significant in that decision: balancing religious and family expectations with societal integration, and the potential impact of veiling on economic participation due to discrimination.

Why do women veil? France's secular policies against veiling have been justified by the idea of a "silent majority" of Muslim women who are forced to wear the veil by their families or communities. According to this idea, the benefits of helping this silent majority outweigh the harm imposed on other female Muslims who truly want to veil (Maurin and Navarrete-Hernandez 2023). However, existing evidence on the motives behind veiling behavior contradicts this argument. In fact, interviews and surveys conducted in France suggest that the vast majority of Muslim women who wear the veil do so by individual choice and not out of coercion (IFOP 2019, Institut Montaigne 2016). Even within the Muslim community, the motives behind veiling seem to be misinterpreted. For instance, non-veiled Muslim women are more likely to believe that veiling is done out of coercion or imitation (IFOP 2019). This discrepancy highlights a key limitation of interview data: it is unclear whether "individual choice" reflects the preferences of the women themselves, or the internalization by these women of the preferences of their social networks.

In a series of interviews with Muslim girls and women,[8] Gaspard and Khosrokhavar (1995) identified three broad categories of veiled women: "veiled immigrants," i.e. middle-aged women

[6] The question of veiling in public resurfaced for instance during the debates surrounding the adoption of the "law on separatisms" of August 2021, with some Senators suggesting a complete ban of all religious symbols in public spaces (see Sénat 2021).

[7] Gresh (2020).

[8] Gaspard and Khosrokhavar (1995) conducted around one hundred interviews with Muslim girls and women in the Paris and Dreux suburbs.

who arrived in France veiled and kept the practice; adolescent girls born in France who wear the veil either by force or by choice; and young women who wear the veil willingly to reconcile their religious duties and integration into French society. The veil worn by first-generation immigrants is well tolerated by French society. Animosity is instead directed towards the veils worn by adolescents and young women born in France, which is perceived as a symbol of failed integration – "a sign of inherent non-Frenchness" (Scott 2009, p. 15).

When asked why they wear the veil, Muslim women mostly invoke religious duty (76%) and issues of safety (35%) (Institut Montaigne 2016). Young women in particular mention "the difficulty to reconcile their families' demands with those of the society" (Khosrokhavar 2004 p. 90). Familial pressures typically discourage them from engaging in activities that favor their integration, such as going out with friends or finding a job. In this respect, veiling can be a tool which allows them to "exempt themselves from the constraints that traditionally weigh on women" (Gaspard and Khosrokhavar 1995, p. 37) and to resolve the tension between religious duty, families' demands, and integration.[9]

This interpretation of veiling as facilitating integration is in line with research in economics which has explored veiling practices in relation to economic participation (Carvalho 2013, Shofia 2020). The theory of Carvalho (2013) considers veiling as a technology available to Muslim women in order to alleviate the intrinsic and social costs of their integration. By providing a practical protection against opportunities to engage in religiously prohibited behaviors, veiling acts both as a commitment to oneself and as a signal of this commitment to others. This commitment aspect of veiling is confirmed by survey evidence and interviews conducted in France and elsewhere.[10] Furthermore, Shofia (2020) provided evidence for this mechanism in a study of veiling among Indonesian schoolgirls.

Veiling and economic participation. The sociological and anthropological record documents the challenges faced by veiled women in France when trying to integrate into the workforce

[9]The following interview excerpts collected by Atasoy (2006) in Canada also illustrate this tension well:

> "It is hard as a young woman not to have a boyfriend in this society. [...] The veil reminds you that this isn't allowed [in Islam]."
>
> Sarah believes the veil keeps her away from doing "stupid things like dating a guy."
>
> "The veil reminds me that I submit to Allah... If I don't wear it, people might take it as I'm doing something wrong."
>
> "If you are not covered, you feel isolated from other Muslim girls. They don't socialize with you. They think you are doing bad things."

[10]See for example Atasoy (2006) for Canada and Read and Bartkowski (2000) and Droogsma (2007) for the United States.

(Adida et al. 2010, 2016, Jouili 2020). Alongside the policies restricting religious expression in public areas, veiled women encounter various constraints in the workplace. For example, French civil servants have an obligation of religious neutrality – a strict application of *laïcité*, the French conception of state secularism. This obligation prohibits the expression of religious beliefs while on duty, including the wearing of conspicuous religious symbols. Breaching this obligation is considered a serious offense that can lead to sanctions or even dismissal.

Veiled women also encounter obstacles in the private sector (Ajbli 2011). Private-sector workers providing a public service are also subject to neutrality requirements. Furthermore, studies have shown that Muslims, particularly those who display higher levels of religiosity (a trait associated with wearing the veil), face discrimination when seeking employment. Valfort (2020) uses a correspondence-test method to demonstrate that while signalling religiosity increases call-back rates for Christian applicants, it significantly reduces them for Muslim applicants in France.[11] Similar discriminatory hiring practices have been reported in other European countries.[12]

Employers claim that discrimination against Muslims is due to religious expression causing conflicts, and accommodating religious practices is viewed as a challenge (Adida et al. 2016, Cintas et al. 2012). Muslims, in particular, face discrimination as some of their religious practices, such as daily prayers and fasting, are perceived as reducing productivity (Bouzar and Bouzar 2009, Maillard 2017).[13] In its yearly surveys of French managers, the *Observatoire du Fait Religieux en Entreprise* documents a rise in observed religious behaviors requiring managerial intervention, with Islam being by far the most cited religion (Institut Montaigne 2014–2021).[14]

Of course, Muslim women report wearing the veil for various other reasons, including signaling piety to potential husbands, or even fashion (Patel 2012). Worth mentioning are identity motives that are not necessarily religious. For some Muslim women, the veil is a means to affirm their distinction with the rest of society and to feel closer to their community of origin (Silhouette-Dercourt et al. 2019). For instance, adolescents who want to distinguish themselves from their

[11] Valfort (2020) uses extra-curricular activities (volunteering for a Christian or a Muslim Scout association) as a signal of religiosity.

[12] Weichselbaumer (2020) and Fernández-Reino et al. (2022) also use correspondence tests to confirm the existence of discrimination against veiled women in Germany, the Netherlands, and Spain.

[13] Hu and Wang (2021) provides empirical evidence suggesting that Ramadan fasting does not in fact reduce productivity.

[14] The *Observatoire du Fait Religieux en Entreprise* conducts surveys on religious behaviors in the workplace. Islam is most frequently associated with observed religious behaviors (73% in 2021), and the proportion of observed religious behaviors requiring managerial intervention has risen from about 25% in 2014 to over 50% in 2021. Of those cases requiring intervention, 19.5% resulted in conflicts in 2021, compared to 6% in 2014. When discriminatory situations in hiring are observed, they involve Muslims in 70% of cases, according to manager reports. In addition, 10% of managers feel overburdened by religious behaviors in their company (Institut Montaigne 2014–2021).

peers may use the veil as a visible sign of difference from the "rooted French" (Khosrokhavar 2004, van der Hasselt 2019). In some cases, wearing the veil is a form of rebellion against a society that claims to defend liberty of choice but discriminates against Muslims, as evidenced by studies on "identity backlash" (Abdelgadir and Fouka 2020).[15]

2.3 Data and descriptives

In this section we start to explore empirically the relationship between veiling behavior and economic participation. We present our main data sources, and we describe them along several dimensions of interest.

2.3.1 Data

Our primary data source is the cross-section from the *Trajectoires et Origines* survey (henceforth TeO; Beauchemin et al. 2016). Conducted in 2008–2009 by the French National Institute for Demographic Studies (INED) and the National Institute of Statistics and Economic Studies (INSEE), the TeO survey targeted adults between 18 and 60 years old residing in metropolitan France. Purposefully oversampling immigrants and minorities, it includes 3,033 women who identify as Muslim. To our knowledge, this is the largest sample of this kind in France.

The TeO dataset is a comprehensive source of information on various aspects of respondents' lives, including living conditions (such as employment, education, housing, commune of residence, and health), social life (such as migration history, language use, family, and children), and public life (such as political views, experiences of discrimination, and social relationships). Of particularly value for this study is the religion section, which is a rare inclusion in French surveys since the collection of individual information on religion is closely monitored. This section includes variables such as religious affiliation, measures of religiosity, religious symbols worn, and intergenerational religious transmission.

We also use the TeO survey to create a panel dataset of respondents' lifetime education and labor-market status. The dataset is constructed by analyzing respondents' retrospective accounts, year by year, of their work status including salaried work, self-employment, unemployment, studying, staying at home, inactive for other reasons, or out of metropolitan France.

Our second data source is the *Annuaire des mosquées de France* (La Boussole 2004), a comprehensive directory of mosques and Muslim praying rooms in France. This is a novel data source

[15] See also Fouka (2020) and Sakalli (2019) for evidence of cultural backlash against assimilation policies in other contexts.

in the literature, which we digitized manually. Compiled by a Muslim association in 2003–2004, the directory provides for each worship facility at the time its full address and estimated capacity by gender.

2.3.2 Measurement

Alongside standard metrics of economic activity, our empirical analysis relies on measures of religious practice and religious social pressure which we describe here.

Veiling. We use the following question from the TeO survey:

> *In your daily life, do you wear in public a piece of clothing or jewelry that might evoke your religion? (1) Never (2) Sometimes (3) Always*

If applicable, respondents were subsequently asked to report which religious symbols they wear. Answers were later sorted by the survey institute into four categories: Jewelry, Clothing, Headcoverings, or Others.

Because they visibly signal religion and are the ones usually targeted by secular policies, we group the Clothing and Headcoverings categories together as *conspicuous symbols*. Among Muslim women this is an excellent proxy for veiling, since headcoverings represent 93% of these conspicuous symbols. In contrast, we group Jewelry and Other symbols, which can usually be hidden, as *discreet symbols*.[16] We then cross these categories with the initial answer on frequency of wearing religious symbols. Thus, in our measure of veiling each respondent is categorized as wearing either (1) *no symbol* (if they answered *Never* to the initial question), (2) *sometimes discreet symbols*, (3) *always discreet symbols*, (4) *sometimes conspicuous symbols*, or (5) *always conspicuous symbols*.[17]

Individual religiosity. The TeO survey includes several questions which relate to individual religiosity. Our preferred measure is the frequency of attendance of religious ceremonies, a standard measure of religiosity which focuses on religious practice (Iyer 2016). To analyze incentives for veiling we combine this measure with other questions related to individual religiosity: the self-reported importance of religion in the respondent's life, whether she uses her

[16] A respondent who wears both discreet and conspicuous symbols is categorized as wearing conspicuous symbols.

[17] A limitation of this data is that appreciations like "sometimes" or "always" remain subjective. For instance, a woman who removes her veil in the workplace by obligation might still consider that she "always" wears it – when she is able to. In our data, a few Muslim women do report veiling "always" even though they work in the public sector, where conspicuous religious symbols are prohibited (cf. section 2.2).

religion to self-identify, the respect of religious dietary restrictions, and religious marriage. In order to aggregate the answers to these questions into a single measure of individual religiosity, we use a measurement system to construct a latent index of individual religiosity, as in Heckman et al. (2013) or Bolt et al. (2021). The advantage of this method is that we are able to leverage the variation on several survey questions while keeping the convenience of a single, continuous measure of religiosity. (In Appendix 2.7.1 we provide details on the procedure and on the survey questions.)

Family and community pressures. As discussed in section 2.2, religious social pressures play a role in women's decisions to integrate socio-economically and to veil. Drawing on insights from the literature on cultural transmission (Bisin and Verdier 2000), and particularly on the distinction between vertical transmission (from parents to children) and horizontal transmission (between peers), our measures of social pressure aim to disentangle the respective influences of women's families and of their larger communities on their decisions.

To capture vertical religious pressure by parents, our preferred measure is a question on the self-reported importance of religion in the respondent's education. We also use whether or not the respondent has a religious first name.[18] As for individual religiosity, we then combine these measures into a single index.

For social pressure stemming from the local community, our preferred measure is the share of Maghrebi immigrants in the neighborhood (IRIS level).[19] We also use a second measure, the local worship capacity per thousand inhabitants for all TeO respondents. We construct this measure using our novel data on Muslim worship facilities in France, by combining information on the place of residence of TeO respondents' with the addresses and estimated capacity of these worship facilities. Since these measures are already continuous, we use them as they are and do not aggregate them into an index.

2.3.3 Descriptive evidence

Using the TeO data, we provide novel summary statistics on Muslim women in France. We provide new empirical evidence for the negative relationship which exists between veiling and

[18] Name-giving has been recognized as an important cultural transmission channel (Fryer and Levitt 2004, Abramitzky et al. 2020, Algan et al. 2022). We classify as religious the names of the Islamic prophet's wives, Khadija, Sawda, Aicha, Hafsa, Zainab, Hind, Juwairiya, Safiya, Ramla, and Maimuna (Morsy 1989); and of his daughter Fatima. Variations in spelling are permitted. For male first names, we follow Sakalli (2019) by considering a name as religious if it is a variation of the prophet's name (Mohamed in French) or if it begins with "Abd-" ("servant of…" in Arabic).

[19] Having a parent (especially a father) born in Maghreb is a strong predictor of Muslim affiliation in France (Abdelgadir and Fouka 2020).

Figure 2.1: Geographical distribution of Muslim women in the TeO survey.

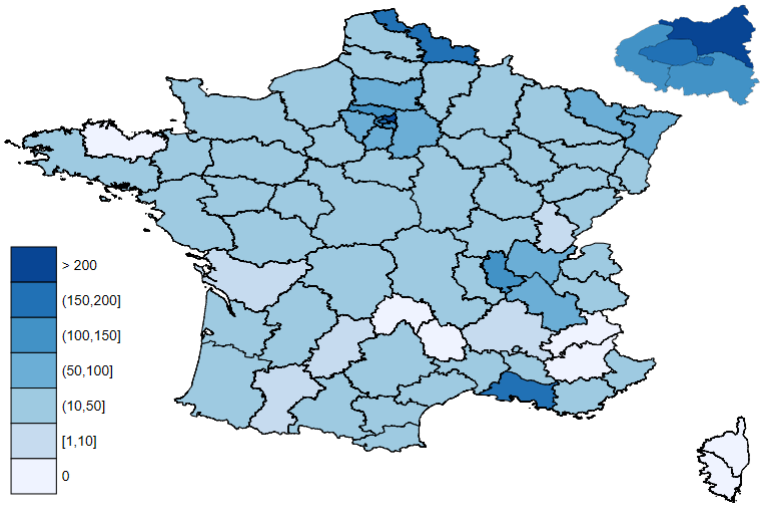

Note: Number of places of residence of Muslim women in the TeO survey per *département*. Some *départements* are collapsed together when counts are low due to confidentiality reasons. The top-right subfigure zooms in on Paris and its suburban area.

economic participation among Muslim women in France, as already suggested by the ethnographic evidence outlined in section 2.2. Through summary statistics and regression analysis, we confirm that veiling is associated with reduced economic participation. Our preferred specification suggests that consistently wearing a conspicuous religious symbol is associated to the same decrease in economic activity as having an additional 1.4 preschool-age children.

Geographical coverage

The representativeness of the ethnographic studies discussed in section 2.2 is limited due to their predominant focus on the Parisian suburbs, some of which are distressed areas that may not accurately reflect the living situations of Muslim women as a whole. In contrast, the TeO survey includes Muslim women from a diverse range of locations, as illustrated in Figure 2.1. Although some respondents remain concentrated in major urban centers such as Paris, Marseille, and Lille, the survey has a wide geographical coverage across the country.

Summary statistics

Table 2.1 presents summary statistics for our main variables of interest, disaggregated by veiling behavior. Panel A examines demographic characteristics and economic outcomes, such as employment and educational attainment. The data reveals that veiled Muslim women have significantly worse economic outcomes compared to those who wear no symbol or discreet ones. On average, they are much less educated, less likely to be employed, and have fewer years of

Table 2.1: Summary statistics by veiling status, Muslim women.

	Veiling behavior				
	No symbol	Sometimes discreet	Always discreet	Sometimes consp.	Always consp.
Panel A: demographics and economic outcomes					
Demographics					
Age in 2008	35.55	28.40	25.06	35.94	36.00
First-gen. immigrant	0.61	0.24	0.51	0.68	0.78
Second-gen. immigrant	0.39	0.66	0.49	0.32	0.22
Number of children	1.78	1.11	0.63	2.26	2.79
Lives in a couple	0.59	0.49	0.48	0.68	0.74
Not a French speaker	0.07	0.02	0.01	0.14	0.32
Labour-force status in 2008					
Employed	0.54	0.43	0.36	0.44	0.22
Unemployed	0.18	0.23	0.27	0.12	0.09
Inactive	0.19	0.15	0.24	0.30	0.65
Student	0.09	0.20	0.13	0.14	0.03
Has never worked	0.19	0.29	0.48	0.31	0.50
Schooling attainment and work experience					
Completed high school	0.78	0.85	0.58	0.68	0.61
Higher education degree	0.22	0.24	0.10	0.20	0.19
Years of schooling	15.30	17.41	15.69	12.86	11.11
Years of work experience	7.06	3.93	3.44	5.75	2.66
Panel B: religious characteristics					
Attends religious ceremonies					
Familial ceremonies only	0.29	0.32	0.29	0.30	0.18
Religious feasts only	0.20	0.34	0.21	0.32	0.27
Once or twice a month	0.03	0.05	0.02	0.08	0.09
At least once a week	0.02	0.01	0.03	0.11	0.19
Importance of religion in education received					
A little important	0.18	0.11	0.05	0.04	0.04
Quite important	0.31	0.30	0.29	0.33	0.15
Very important	0.47	0.58	0.64	0.63	0.81
Percentage of Maghrebi immigrants in neighborhood					
Fourth quintile	0.27	0.31	0.19	0.32	0.32
Top quintile	0.43	0.40	0.54	0.44	0.47
Observations	2,017	166	151	148	516

Note: This table reports means of variables of interest by veiling status as defined by the type of symbol and the frequency at which they are worn. Observations are weighted using the survey weights provided in TeO dataset.

work experience, despite being older. Particularly striking is the sharp difference in activity rates (activity being defined as either working, looking for a job, or studying). Almost two-thirds of women who always veil are inactive, compared to less than 20% for non-veiled women, indicating significant barriers to integration linked to the veil.

Panel B examines our primary measures of religiosity and religious social pressure. We observe a positive link between both individual religiosity and veiling, and religious social pressure. On average, veiled Muslim women attend religious ceremonies more frequently, received an education which stressed the importance of religion more, and they now live in neighborhoods with higher proportions of Maghrebi immigrants. Our other measures of religiosity and religious social pressure confirm these patterns (Appendix Table 2.6).

Veiling is negatively correlated with economic participation

Our summary statistics provide some preliminary evidence of the negative link between veiling and economic participation, which we now investigate further using regression analysis. We perform two regression exercises, which complement each other.

First, we explore the relationship between Muslim women's active status and veiling in the cross-section. With this approach, we are able to include a rich set of controls by using the wide range of information on respondents available in the TeO survey. We also check the robustness of our results by restricting attention to particular subsamples and by conducting placebo tests on populations other than Muslim women.

Our second approach is to explore this relationship in the panel dataset that we constructed from respondents' retrospective accounts of their studies and professional trajectories. Since this retrospective account does not include most other questions in the survey, our set of controls is restricted in this approach. However, the panel dimension allows us to verify that the relationship between veiling and economic activity status is not merely due to the particular timing of the survey. Timing might indeed be a concern since the survey was conducted around the time of the Great Recession, which may have affected veiled women disproportionately, e.g. if they face stronger discrimination. Together, the two exercises thus provide a robust assessment of the correlation between veiling and economic participation.

Cross-sectional analysis. Table 2.2 shows the results of linear regressions where the outcome variable is the activity status (0 if inactive, 1 if active), and the main explanatory variable is the respondent's veiling behavior. Other important explanatory variables include our measures of individual religiosity and religious social pressure, economic characteristics which are

Table 2.2: Veiling and economic participation, Muslim women.

	Woman is active (= 1 if active, = 0 if inactive)					
	(1)	(2)	(3)	(4)	(5)	(6)
Veiling behavior						
Sometimes discreet symbol	0.029	-0.052	-0.054	-0.054	-0.028	-0.036
	(0.041)	(0.037)	(0.035)	(0.037)	(0.036)	(0.034)
Always discreet symbol	0.117***	0.019	0.036	0.028	0.038	0.055*
	(0.028)	(0.028)	(0.029)	(0.028)	(0.029)	(0.031)
Sometimes conspicuous symbol	-0.107*	-0.090**	-0.072	-0.083*	-0.053	-0.055
	(0.055)	(0.046)	(0.044)	(0.047)	(0.038)	(0.037)
Always conspicuous symbol	-0.441***	-0.351***	-0.300***	-0.328***	-0.268***	-0.232***
	(0.048)	(0.032)	(0.030)	(0.033)	(0.031)	(0.030)
Demographics						
Number of children			-0.051***			-0.028**
			(0.012)			(0.011)
Number of children below age 4			-0.165***			-0.171***
			(0.022)			(0.020)
Lives in a couple			-0.065**			-0.065*
			(0.031)			(0.033)
Educational attainment and work experience						
Years of schooling					0.011***	0.010***
					(0.003)	(0.002)
Completed high school					-0.017	-0.021
					(0.025)	(0.024)
Higher education degree					0.059**	0.040*
					(0.024)	(0.021)
Years of work experience					0.041***	0.039***
					(0.004)	(0.004)
Experience squared					-0.001***	-0.001***
					(0.000)	(0.000)
Constant	0.812***	0.514***	0.341**	0.481***	0.055	-0.144
	(0.016)	(0.138)	(0.140)	(0.157)	(0.146)	(0.157)
Other demographic controls			✓			✓
Religiosity controls				✓		✓
Religious influence controls				✓		✓
Birthyear dummies		✓	✓	✓	✓	✓
Age of arrival in France dummies		✓	✓	✓	✓	✓
Birthplace dummies		✓	✓	✓	✓	✓
Region of residence dummies		✓	✓	✓	✓	✓
Observations	2433	2433	2433	2433	2433	2433
R^2	0.147	0.358	0.428	0.374	0.450	0.511

Note: This table reports results of linear regressions on a dichotomous variable taking the value of 1 if a woman reports being in the labor force or studying. The sample is restricted to Muslim women with no missing covariates. Observations are weighted using the weights provided in the TeO survey. Robust standard errors in parentheses. Level of statistical significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

usual predictors of labor market participation such as education and experience, and other demographic predictors. The sample is restricted to Muslim women with non-missing covariates, yielding 2433 observations.

Column (1) includes veiling behavior as the only predictors of active status. Veiling behavior alone is an important predictor of the activity rate, explaining 13.5% of the variation in the activity status. In columns (2) to (6) we add more controls, including dummy variables for birth year, age of arrival in France, birthplace, and region of residence. We further include a set of dummy variables capturing the conditions in which the survey took place (whether the partner was present, whether parents were present, survey month dummies, age group of surveyor dummies, and surveyor's gender), which gives us confidence that social desirability bias is minimized in our regressions.

We include additional groups of control variables one by one to investigate the relative contribution of different mechanisms. The last column reports the results of a regression controlling for all of the covariates. In this last specification, the only significant predictors of the activity status are the wearing of conspicuous symbols, the number of children, age, birthplace, and the education level. The magnitude of the main coefficients of interest is reduced compared to specifications with a sparser choice of controls, but it remains statistically and economically significant. The point estimates indicate that Muslim women who always wear a conspicuous symbol are 23 p.p. less likely to be active compared to those who never wear any symbol. Even in this most parsimonious specification, the estimated effect is substantial: it is equivalent to having an additional 1.4 preschool-age children.

Robustness checks. Overall, the regression results of Table 2.2 confirm a strong negative association between veiling and economic participation. We further verify the validity of this statement through a series of robustness checks, the results of which are summarized in Table 2.3. The first three columns correspond to re-estimations of our preferred specification (column 6, Table 2.2) in different subsamples. The goal of this exercise is to verify that our results are not driven by particular observations or simply capturing something else apart from the potential impact of veiling. The first row excludes students to use a more conventional measure of economic participation, that is, labor-market participation. The second row excludes individuals born outside France, since summary statistics suggested an important difference in immigration status between veiled and non-veiled women. The third row excludes women whose religious symbol is categorized as Other (i.e. neither Clothing, Headcoverings, or Jewelry). Restricting attention to these subsamples yields point estimates for the effect of conspicuous-symbol wearing

Table 2.3: Robustness checks, cross-sectional data

				Other religious groups (placebo)		
	Excl. students (1)	Born in France (2)	Excl. "other" symbols (3)	Muslim men (4)	Excl. Muslims and Catholics (5)	All non-Muslims (6)
Veiling status						
Sometimes discreet	-0.033	0.035	-0.040	0.021	0.017	-0.017
	(0.040)	(0.032)	(0.034)	(0.013)	(0.018)	(0.012)
Always discreet	0.072	0.049	0.058	0.040***	-0.008	-0.022*
	(0.037)	(0.034)	(0.031)	(0.012)	(0.019)	(0.012)
Sometimes conspicuous	-0.063	0.075*	-0.058	-0.050*	0.022	0.012
	(0.044)	(0.036)	(0.037)	(0.029)	(0.068)	(0.044)
Always conspicuous	-0.234***	-0.246***	-0.226***	0.016	0.080	0.066
	(0.031)	(0.052)	(0.030)	(0.078)	(0.154)	(0.141)
Controls	✓	✓	✓	✓	✓	✓
Observations	2,158	1,199	2,427	2,197	1,756	5,744
R^2	0.510	0.411	0.517	0.204	0.245	0.196

Controls included in the regressions are the full set of variables included in Table 2.2, column (6). In column (1), we exclude students so that the dependent variable becomes labor-market participation. In column (2), the estimation sample is restricted to second-generation immigrant Muslim women (born in France of foreign parents). In column (3), individuals reporting to wear a religious symbol that is neither jewelry, a headcovering, or clothing (symbols labelled as "other") are excluded from the sample. Columns (4) to (6) estimate the same regression on other religious groups. Robust standard errors in parentheses. Level of statistical significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

which are of similar magnitude to those obtained on the complete sample.

Columns four to six of Table 2.3 re-estimate the same specification, this time on groups other than Muslim women, thus providing a form of placebo test. We find that wearing a religious symbol has no significant association with economic participation for Muslim men, nor for women and men with different religious affiliations. These results confirm the unique place of the Islamic veil among other religious symbols, as evidenced by the debates mentioned in section 2.2. Whether it is because of individual preferences, social pressure, legal restrictions on veiling at work, or discrimination, veiling seems to be the only widespread religious symbol which is strongly associated with decreased economic participation.

Panel analysis. We perform another robustness check in order to control for timing effects, in particular in the event that veiled women's employment prospects were differentially affected by the 2008 economic crisis (which coincided with the time of the survey). To investigate this possibility, we use the retrospective panel dataset, where we exclude observations for which individuals report multiple activities as well as periods in which the respondent was out of metropolitan France. This empirical strategy allows us to control for time-varying observables

and time fixed effects, to substantially increase the number of observations, as well as to include random effects. For the sake of space, we present this analysis in Appendix 2.7.1. The results overall confirm the findings obtained in the cross-sectional analysis, with the wearing of a conspicuous symbol being associated with a significant decline in economic participation, similar in magnitude to that obtained in the cross-sectional analysis.

2.4 A model of veiling and labor supply

In the previous section, we have shown that veiling displays a strong negative association with economic participation in France. Our discussion of the literature on veiling from section 2.2 suggests that such an association can originate from two sources of incentives, namely religious (women who veil are more religious and therefore less likely to engage with an environment they perceive as dangerous) and economic (women who veil face discrimination on the labor market).

In order to structure our empirical analysis of these motives, in this section we model Muslim women's joint decision of economic participation and veiling. This model notably builds on the theory of Carvalho (2013), who considers the veil as a tool available for women to mitigate the socio-religious cost of their integration. We expand on this model by proposing a general analytical framework which remains agnostic as to the reasons why women veil. We then show that this general framework can be specified to accommodate together both religious motives in the spirit of Carvalho (2013), as well as economic motives stemming from anti-veil discrimination on the labor market.

2.4.1 General model

We consider a static model in which an agent must simultaneously decide on her labor supply and her veiling behavior. For her labor supply, she allocates her total time budget $T = 1$ between time worked, t, and time devoted to leisure, $1-t$. In addition, she chooses what degree of veiling to adopt at work, v_1, and what degree of veiling to adopt during her leisure time, v_0. The flow utility that the agent derives from work and leisure then depends on her degree of veiling in each of these activities. The model remains agnostic about whether veiling has a positive or negative effect on the flow utility of working or leisure. In this way, it is able to account for a wide range of mechanisms linking veiling behaviors and labor supply decisions, from the religious stigma faced by working Muslim women to identity-based discriminations at and outside work.

Formally, the utility that the agent derives from working, u_1, and the utility that she derives

from leisure, u_0, take the form

$$u_j(v_j) = a_j + b_j\, v_j - c(v_j). \tag{2.1}$$

The parameters a_j and b_j are constants which are specific to activity j. They may be positive or negative. The parameter a_j represents the baseline return to activity j when not veiling. It could account for motivations as diverse as, for instance, the agent's baseline wage rate, the religious social pressure that she might face against her working, or how much she appreciates her colleagues.

The parameter b_j represents how veiling affects this baseline return to activity j. This could be a combination of positive effects, such as alleviating the religious stigma faced by working Muslim women (as in Carvalho 2013); and negative ones, such as triggering discriminations or hostile reactions from peers.

Finally, there is an intrinsic cost $c(\cdot)$ to wearing the veil, which is the same across activities j. Following Carvalho (2013), this cost can for instance be interpreted as physical discomfort. We assume that the cost function $c(\cdot)$ is convex, with $c'(0) = 0$ and $\lim_{v\to 1} c'(v) = \infty$.

With this, we can write the complete utility function of the agent:

$$\begin{aligned} U(t, v_1, v_0) &= t\, u_1 + (1-t)\, u_0 - d(t) \\ &= t\, \big[\underbrace{a_1 + b_1\, v_1 - c(v_1)}_{\text{flow utility from work}}\big] + (1-t)\, \big[\underbrace{a_0 + b_0\, v_0 - c(v_0)}_{\text{flow utility from leisure}}\big] - d(t). \end{aligned} \tag{2.2}$$

The component $d(t)$ represents a disutility of working, which we assume is increasing and convex, with $d'(0) = 0$ and $\lim_{t\to 1} d'(t) = \infty$.

2.4.2 Optimal choices

The problem of the agent is to find the time allocation t and the degrees of veiling v_0 and v_1 which maximize her utility (2.2). This problem can be solved sequentially. First, the agent determines for each activity the degree of veiling which maximizes her flow utility. Second, she chooses her labor supply based on those optimized flow utilities.

Veiling. Call v_j^* the optimal degree of veiling in activity j. If the agent has negative returns to veiling in activity j, i.e. $b_j \leq 0$, then she has no incentive to veil in that activity and her optimal degree of veiling is $v_j^* = 0$. Otherwise, if $b_j > 0$, then her optimal degree of veiling maximizes

the utility (2.1) that she derives from activity j:

$$c'(v_j^*) = b_j, \tag{2.3}$$

so that v_j^* is positive, increasing in the agent's return to veiling b_j.

Thus, in this model, differences in veiling behavior between work and leisure time are reflective of different returns to veiling for the agent across these activities. We summarize this result in the following lemma, which will become useful later on.

Lemma 2.4.2.1. *The agent veils more at work than during leisure time if and only if* $b_1 > b_0$.

We now move on to the labor supply problem. In what follows, we denote by $u_j^* = u_j(v_j^*)$ the indirect utility that the agent derives from activity j.

Labor supply. Call t^* the optimal labor supply. If her indirect utility obtained from working is less than that obtained from her leisure, i.e. $u_1^* \leq u_0^*$, then the agent has no incentive to work and her optimal labor supply is $t^* = 0$. Otherwise, if $u_1^* > u_0^*$, her optimal labor supply t^* solves the first-order condition

$$d'(t^*) = u_1^* - u_0^*, \tag{2.4}$$

so that t^* is positive, increasing in the indirect utility of working u_1^*, and decreasing in the indirect utility of leisure u_0^*.

In the case whereby the agent has equal returns on veiling for both activities, i.e. $b_1 = b_0$, veiling has no impact on the labor supply decision. Indeed, in this case the agent chooses the same degree of veiling at work and during leisure time: $v_1^* = v_0^*$. Therefore the difference in indirect utilities is simply $u_1^* - u_0^* = a_1 - a_0$, which depends only on the baseline return to each activity. Thus, veiling has an impact on the agent's labor supply decision only if it distorts the returns to work and leisure in distinct ways. In particular, if veiling motives are purely personal and do not interact with the environment, the veiling and labor supply decisions are orthogonal.

2.4.3 Mechanisms

In this section we provide two concrete examples of theoretical mechanisms which may underpin the relationship between the veiling and labor supply decisions, based on the discussion from section 2.2. We use these examples to provide micro-foundations to the generic parameters a_j and b_j that we have introduced in our general framework above.

We begin by examining the theoretical mechanism studied by Carvalho (2013), which relates to social norms and expectations. In some communities, women may face social pressure to limit

their labor supply in order to conform to gender role expectations and maintain social approval. This social pressure can be amplified for religious women who may themselves feel reluctant to integrate into a work environment they perceive as religiously unsafe. Here, veiling can serve a dual purpose as a self-commitment to religious beliefs and as a signal to their community of their religious intentions. As a result, veiling can help mitigate the social cost of women's employment, making it a useful tool for their economic integration.

Second, we consider a mechanism which relates to the role of discrimination. Veiled women may face discrimination in the workplace due to negative stereotypes or biases held by their employers or colleagues. This discrimination may limit their opportunities for employment or career advancement, and could ultimately lead them to reduce their labor supply. We predict that women with higher wage potential, who face a greater opportunity cost of unemployment or limited career advancement, will incur higher costs associated with veiling.

Religious motives: the Carvalho model. Let us show that the Carvalho model of veiling is a particular case of the framework that we have developed above. In the Carvalho model, the incentive to veil stems from a combination of the individual religiosity of the agent, r, and of the religious social pressure, R. Together, these religious factors determine the penalty that the agent suffers if she engages in religiously-prohibited behavior. This penalty, equal to $-(r+R)$, is both self- and socially-imposed, reflecting personal regret on the one hand, and social stigma on the other hand. It is steeper if the agent herself has higher religiosity, and if there is more religious social pressure. Note that in this context both r and R can be negative, meaning individual or social approval for religiously-prohibited behavior.

Each activity j, working or leisure,[20] is then characterized by an exogenous risk of engaging in religiously-prohibited behavior, p_j. Crucially, the agent is able to attenuate that risk by veiling. Specifically, if she chooses a degree of veiling v_j, then the probability that she engages in religiously-prohibited behavior becomes $p_j(1-v_j)$. Veiling also entails a cost $c(v_j)$ (e.g. physical discomfort).

Finally, there is a material reward m_j associated with each activity j. As a result, the expected utility that the agent derives from activity j is

$$u_j(v_j) = -p_j(1-v_j)(r+R) - c(v_j) + m_j. \tag{2.5}$$

This utility function is a particular case of equation (2.1), which is obtained by taking $a_j = -p_j(r+R) + m_j$ and $b_j = p_j(r+R)$.

[20] Carvalho gives a broader interpretation of this decision as a choice between *integration* or *segregation*.

In the Carvalho model, the exogenous risk of engaging in religiously-prohibited behavior is assumed to be greater at work than during leisure time: $p_1 > p_0$. This assumption implies that a woman will always choose a higher degree of veiling at work than during leisure time. Indeed, recall that for the agent to veil at all, she must have positive returns to veiling, i.e. $b_j > 0$. For this to hold here, the agent must have $r + R > 0$, and as a consequence $p_1(r + R) > p_0(r + R)$, i.e. $b_1 > b_0$. Thus, according to our lemma 2.4.2.1, a woman will always veil more at work than during leisure time in the Carvalho model.

Regarding the choice of activity, Carvalho considers a discrete choice $j \in \{0, 1\}$. Again this is a particular case of our framework, obtained by ignoring the disutility of working: $d(t) = 0$. Following our analysis of the labor supply decision, the agent will work if her indirect utility from working is greater than that from leisure, $u_1^* > u_0^*$. This happens if and only if the material reward for working m_1 is large enough.

Here Carvalho shows an interesting result, namely that within a range of values of this material reward m_1, (i) low-religiosity women choose to work, (ii) high-religiosity women choose not to work, and (iii) low-religiosity women veil more than high-religiosity ones. This happens provided that the surrounding population approves of the veil, i.e. $R > 0$, because in this case low-religiosity working women choose to attenuate the social penalty associated with working by veiling. Shofia (2020) finds evidence for this pattern of veiling among women in Indonesia.

Economic motives: labor market discrimination against veiling. Our general model above can also account for discrimination against veiling on the labor market. Consider a simple consumption–leisure framework: the agent has quasilinear utility $U(x, t) = x + g(1 - t)$ where x is her consumption of a numeraire good and $1 - t$ is her leisure (the function $g(\cdot)$ is increasing and concave). Consumption is the only source of spending, so that the budget constraint is $x = w\,t$, where w is the agent's wage rate.

We assume that discrimination against veiling has a direct negative effect on the agent's effective wage. Indeed, such discrimination can typically make it more difficult for women who wear the veil to secure and keep a job or to advance in their career (cf. section 2.2). This suggests that the financial cost of discrimination may be greater for women with higher earning potential. For example, the opportunity cost of job loss or slower career progression is proportional to one's earning potential.[21] Therefore, we assume that for an agent with wage potential w, the financial cost of adopting the veiling level v is equal to $-wv$. The budget constraint of the agent can then

[21] That the financial cost of discrimination is higher for higher-paid women does not mean in general that the welfare (utility-related) cost is higher for them (although with a quasilinear utility function this is the case).

be expressed as $x = w(1-v)t$.

Aside from the financial cost of discriminations, suppose that veiling at work provides a return y to the agent (maybe through the religious incentive mechanism discussed above), and entails a cost $c(v)$. In this case, her utility function is

$$U(t,v) = \big[w + (y-w)v - c(v)\big]\, t - d(t) \tag{2.6}$$

where $d(t) = -g(1-t)$, so that the function $d(\cdot)$ is increasing and convex. Again, this is a particular case of the utility function (2.1), obtained by taking $a_1 = w$, $b_1 = y - w$, and $a_0 = b_0 = 0$. This model predicts that women with a higher wage potential w should work more and veil less than those with a lower wage. This result is a direct consequence of veiling having a negative, proportional impact on the agent's effective wage.

The two mechanisms above mostly play in opposite directions. According to the first mechanism, women who are religious or who face religious pressure from their family or community have an incentive to veil at work in order to mitigate the social penalty associated with working. But according to the second mechanism, discrimination at work provides an opposite incentive to unveil at the workplace. In the next section we pool these two motives together in a unified empirical model. We then use data on veiling behaviors and employment of Muslim women to quantify the various effects at hand.

2.5 Empirical analysis

2.5.1 Econometric model

Our econometric specification is derived by pooling together the two motives for (un)veiling described in the previous section, religious and economic. To capture these motives, we focus on three main individual characteristics: individual religiosity r_i, the religious social pressure faced by the individual R_i, and earning potential w_i. We obtain a unified expression for the utility that woman i receives by jointly choosing the degree of veiling v and the activity j:

$$u_{ij}(v) = \underbrace{-p_j(1-v)(r_i + R_i)}_{\text{religious motives}} + \underbrace{\mathbf{1}_{\{j=1\}} w_i(1-v)}_{\text{economic motives}} - c(v). \tag{2.7}$$

Our empirical approach relies on measures of the individual characteristics r_i, R_i, and w_i. We use the data and constructed measures that we described in section 2.3.2. Regarding individual religiosity, we use our index measure aggregated from six different survey questions,

$\mathsf{Religiosity}_i$. Regarding religious social pressure, we use our index measure of vertical pressure, $\mathsf{VertiReligiousPressure}_i$, and two measures of horizontal pressure, $\mathsf{ShareMaghrebi}_i$ (the share of Maghrebi immigrants in the individual's neighborhood) and $\mathsf{MosqueCapacity}_i$ (the local capacity for Muslim worship). Regarding the earning potential, we use measures of both the individual's educational attainment using her years of schooling, $\mathsf{Education}_i$, and her years of professional experience, $\mathsf{Experience}_i$. To summarize, we use the following proxies for the individual characteristic of woman i:

$$r_i \sim \mathsf{Religiosity}_i \tag{2.8}$$

$$R_i \sim \mathsf{VertiReligiousPressure}_i + \mathsf{ShareMaghrebi}_i + \mathsf{MosqueCapacity}_i \tag{2.9}$$

$$w_i \sim \mathsf{Education}_i + \mathsf{Experience}_i. \tag{2.10}$$

Next, we formulate an econometric model informed by the theory which is based on these variables. We use a multinomial logit model to explain the joint decision of activity and veiling, (j, v), with three levels of veiling $v \in \{0, 1, 2\}$, and two activity statuses $j \in \{0, 1\}$. Adapting equation (2.7) into an econometric model which uses the proxies described above, the utility for woman i to jointly choose activity j and veiling level v is given by

$$\begin{aligned} u_{ijv} = \alpha_{jv} &+ \beta^1_{jv} \times \mathsf{Religiosity}_i + \beta^2_{jv} \times \mathsf{VertiReligiousPressure}_i \\ &+ \beta^3_{jv} \times \mathsf{ShareMaghrebi}_i + \beta^4_{jv} \times \mathsf{MosqueCapacity}_i \\ &+ \gamma^1_{jv} \times \mathsf{Education}_i + \gamma^2_{jv} \times \mathsf{Experience}_i + X_i' \theta_{jv} + \varepsilon_{ijv}. \end{aligned} \tag{2.11}$$

Here X_i is a set of individual-level controls, and ε_{ijv} is the unobserved part of the utility. The coefficients β_{jv}, γ_{jv} and θ_{jv} are estimated with respect to a baseline, $(j, v) = (0, 0)$. We assume that the unobserved components of utility ε_{ijv} are distributed i.i.d. Gumbel, giving rise to a standard multinomial logit model in which the probability to choose alternative (j, v) is given by

$$\mathrm{Prob}_{jv} = \frac{\exp u_{jv}}{\sum_{j'v'} \exp u_{j'v'}}. \tag{2.12}$$

2.5.2 Implications of the model

According to the model, the religious motives and the economic motives channels have clear implications on the values of parameters to estimate. Table 2.4 outlines the correspondence between parameters of our estimating equation (2.11) and the theoretical components of the model.

Table 2.4: Correspondence between estimated parameters and theoretical model

Explanatory variable	Parameter	Proportional to...	Varies with v	Varies with j
Religiosity variables				
$\mathsf{Religiosity}_i$	β^1_{jv}	$-p_j(1-v)$	$+$	$-$
$\mathsf{VertiReligiousPressure}_i$	β^2_{jv}	$-p_j(1-v)$	$+$	$-$
$\mathsf{ShareMaghrebi}_i$	β^3_{jv}	$-p_j(1-v)$	$+$	$-$
$\mathsf{MosqueCapacity}_i$	β^4_{jv}	$-p_j(1-v)$	$+$	$-$
Economic variables				
$\mathsf{Education}_i$	γ^1_{jv}	$\mathbf{1}_{\{j=1\}}(1-v)$	$-$	$+$
$\mathsf{Experience}_i$	γ^2_{jv}	$\mathbf{1}_{\{j=1\}}(1-v)$	$-$	$+$

The correspondence of Table 2.4 allows us to establish model implications for our empirical parameter estimates. To lay out these implications, in the following we separate our explanatory variables into two categories. The first category of variables are "religiosity variables" – they correspond to the religious motives behind the joint decision of economic participation and veiling. These religiosity variables are associated with the β_{jv} parameters: β^1_{jv}, β^2_{jv}, β^3_{jv}, and β^4_{jv}. The second category of variables are "economic variables," and they correspond to economic motives. They are associated with the γ_{jv} parameters: γ^1_{jv} and γ^2_{jv}.

We describe below the empirical implications of the religious and economic motives of the model for our parameter estimates. Note that the same implications apply to β^1_{jv}, β^2_{jv}, β^3_{jv}, and β^4_{jv} on the one hand; and to γ^1_{jv} and γ^2_{jv} on the other hand. Therefore, we drop the superscripts 1, 2, 3 and 4 in the statements below and refer to generic parameters β_{jv} and γ_{jv} instead.

Implication 1. Within activity,

(a) religiosity variables have a milder (negative) impact on utility for women who veil more:

$$\text{at } j \text{ fixed}, \quad \beta_{j0} < \beta_{j1} < \beta_{j2},$$

(b) economic variables have a milder (positive) impact on utility for women who veil more:

$$\text{at } j \text{ fixed}, \quad \gamma_{j0} > \gamma_{j1} > \gamma_{j2}.$$

Implication 2. For a given degree of veiling,

(a) religiosity variables have a stronger (negative) impact on utility for women who participate economically:

$$\text{at } v \text{ fixed}, \quad \beta_{0v} > \beta_{1v},$$

(b) economic variables have a stronger (positive) impact on utility for women who participate

economically:

$$\text{at } v \text{ fixed}, \quad \gamma_{0v} < \gamma_{1v}.$$

To interpret these implications of the model, let us focus on the meaning of the parameters to estimate. For instance, the parameter β^1_{jv} indicates how own religiosity impacts the probability of choosing the alternative (j, v). According to the theory, this impact is negative since religiosity implies more limitations on acceptable behavior and a higher intensity of regret. In magnitude, the impact should be milder for women who veil – this is the purpose of veiling in the Carvalho model – hence β^1_{jv} should be increasing in v (Implication 1a). Furthermore, the impact should be greater for working women – because the work environment is more risky than the home environment – hence β^1_{jv} should be decreasing in j (Implication 2a). Similar predictions apply for β^2_{jv}, β^3_{jv} and β^4_{jv}, which relate to the social religious pressure.

Next, the parameter γ^1_{jv} indicates how education impacts the probability of choosing the alternative (j, v). In the model education plays a role by increasing the working wage. Therefore the impact of education should be lower for women who veil more – they have lower expected wage because of discrimination (Implication 1b); and it should be greater for women who work compared to those who do not (Implication 2b). Similar predictions apply to γ^2_{jv}, which relates to professional experience.

Implications 1 and 2 above focus on veiling and economic participation choices independently. However, our main interest is to understand how veiling and economic participation choices interact, and what are the relevant mechanisms in this interaction. In the model there are two such mechanisms: the religious motives channel, inspired by the Carvalho model; and the economic discrimination channel. Because they relate to the interaction between veiling and economic participation decisions, these mechanisms will be captured by studying the signs of double differences in the parameters β_{jv} and γ_{jv}.

According to the religious motives mechanism, the religious benefits of veiling are greater for women who integrate economically. This is stated formally as follows:

Implication 3: Religious motives channel. The religious returns on utility to increasing one's degree of veiling are larger for women who participate economically, compared to those who don't:

$$\text{for } v < v' \text{ fixed}, \quad \beta_{1v'} - \beta_{1v} > \beta_{0v'} - \beta_{0v}.$$

Finally, according to the economic discrimination mechanism, the economic losses induced by veiling are greater for women who integrate economically. This is stated formally as follows:

Implication 4: Economic discrimination channel. The economic returns to being economically active are smaller for women who veil, compared to those who don't:

$$\text{for } v < v' \text{ fixed}, \quad \gamma_{1v'} - \gamma_{0v'} < \gamma_{1v} - \gamma_{0v}.$$

Having established these empirical implications of the model's different mechanisms, we now turn to the estimation and to testing the model implications 1–4.

2.5.3 Results

Table 2.5 presents the results for the estimation of equation (2.11). Recall that all parameter estimates are relative to the baseline of an inactive woman who never wears religious symbols. This estimation is performed without controls – in Appendix 2.7.1 we perform the same exercise while including controls, and observe that results remain sensibly similar.

The parameter estimates suggest two main findings. To ease interpretation, we focus on the predicted marginal effects (panel B in Table 2.5). First, individual religiosity is a strong and significant predictor of changes in veiling behavior, but the same observation does not hold for

Table 2.5: Determinants of joint employment and veiling decision, multinomial logit.

Activity choice (j)	Inactive ($j=0$)			Active ($j=1$)		
Veiling choice (v)	None (*baseline*)	Discreet (1)	Conspicuous (2)	None (3)	Discreet (4)	Conspicuous (5)
Panel A: Parameter estimates						
Indiv. religiosity (β^1_{jv})	0	**0.78** (0.21)	**2.26** (0.28)	0.17 (0.16)	**1.00** (0.21)	**2.18** (0.34)
Vert. pressure (β^2_{jv})	0	–1.74 (2.97)	1.54$^+$ (0.81)	0.05 (0.69)	0.56 (0.88)	0.96 (0.94)
Horiz. pressure						
ShareMaghrebi$_i$ (β^3_{jv})	0	4.14 (3.23)	0.68 (1.20)	0.18 (0.85)	0.28 (1.01)	2.14 (1.39)
CapacityMosques$_i$ (β^4_{jv})	0	–0.15 (0.12)	**0.10** (0.04)	0.02 (0.02)	–0.03 (0.02)	0.05 (0.03)
Schooling (γ^1_{jv})	0	0.03$^+$ (0.02)	–0.03$^+$ (0.02)	**0.15** (0.02)	**0.15** (0.02)	**0.06** (0.02)
Work experience (γ^2_{jv})	0	–0.11$^+$ (0.07)	–0.06* (0.03)	**0.12** (0.02)	**0.08** (0.02)	**0.06** (0.02)
Panel B: Average marginal effects						
Indiv. religiosity (β^1_{jv})	**–0.09** (0.01)	0.00 (0.01)	**0.13** (0.02)	**–0.17** (0.02)	**0.05** (0.02)	**0.08** (0.02)
Vert. pressure (β^2_{jv})	–0.03 (0.07)	–0.08 (0.12)	0.11* (0.05)	–0.08 (0.11)	0.04 (0.08)	0.03 (0.04)
Horiz. pressure						
ShareMaghrebi$_i$ (β^3_{jv})	–0.08 (0.09)	0.13 (0.14)	–0.00 (0.08)	–0.12 (0.14)	–0.03 (0.09)	0.09 (0.07)
CapacityMosques$_i$ ($\beta^4_{jv} \times 10$)	–0.02 (0.02)	–0.06 (0.05)	**0.08** (0.03)	0.03 (0.04)	–0.05* (0.02)	0.02 (0.01)
Schooling ($\gamma^1_{jv} \times 10$)	**–0.10** (0.01)	**–0.02** (0.01)	**–0.09** (0.01)	**0.18** (0.02)	**0.05** (0.01)	–0.02$^+$ (0.01)
Work experience ($\gamma^2_{jv} \times 10$)	**–0.07** (0.02)	–0.06$^+$ (0.04)	**–0.10** (0.02)	**0.21** (0.02)	0.01 (0.01)	0.01 (0.01)
Observations	2802					
Sampling weights	✓					
Pseudo R^2	0.159					

Note: This table reports estimates of the parameters of the econometric model (2.11). The baseline category is the choice of inactivity and not wearing any religious symbol. Individual religiosity and vertical religious pressures are measured as indices (with mean zero and variance 1) constructed from multiple proxies available in the TeO data (see Appendix 2.7.1 for details). ShareMaghrebi$_i$ is the proportion of the local population that is of Maghrebi origin. CapacityMosques$_i$ is the estimated capacity in Muslim places of worship in the area of residence. Robust standard errors in parentheses. Point estimates in bold are significant at the 1% level ($p < 0.01$), * $p < 0.05$, $^+$ $p < 0.1$.

social pressures. For example, we estimate that a 1 standard deviation increase in individual religiosity decreases the probability of not wearing any religious symbol and being active (resp. inactive) by 17 percentage points (resp. 9 p.p.). On the contrary, it increases the probability of wearing a conspicuous symbol and being active (resp. inactive) by 8 percentage points (resp. 13 p.p.). Social religious pressure (both vertical and horiztonal) is also associated with higher degrees of veiling, although most parameter estimates are not significantly different from 0 at the conventional levels. For instance, a 1 s.d. increase in vertical social pressure is associated with an 11 p.p. increase in the probability of wearing a conspicuous symbol and being inactive, while an extra 10 Muslim worship seats per 1000 inhabitants is associated with an 8 p.p. increase in the same probability. Overall, both the magnitude of the estimates and their significance level suggest that individual religious motives are the strongest predictors of veiling behavior, above (and conditional on) other social religious pressures.

Second, both schooling and work experience substantially increase the probability of being active and decrease the probability of veiling. For instance, 10 additional school years are associated with an 18 p.p. increase (resp. 5 p.p.) in the probability of being active and wearing no symbol (resp. wearing a discrete symbol). Interestingly however, these human capital factors are not associated with an increase in the probability of being active while wearing a conspicuous symbol. This result might suggest that veiling at work offsets the benefits of human capital on economic activity, an expected consequence of the labor-market discrimination channel.

We illustrate these results in Figure 2.2 by plotting the utility obtained by veiling for an 'average' woman in our sample, according to our estimates.[22] We observe that this average woman has a disincentive to veil overall if she is active, which is a consequence of the economic motives being stronger than the religious ones. On the contrary, an inactive woman has an incentive to veil, because she is less affected by economic motives.

We then compute the same utilities in a counterfactual, Muslim-majority environment in which there is no economic discrimination against wearing the veil at work.[23] In this case, we see that active and inactive women have somewhat equivalent incentives to veil, which sharply contrasts with our findings in the French setting. Active women benefit slightly more from veiling than inactive ones overall, a finding which is consistent with the religious channel of the Carvalho

[22]We set the following values for this 'average' Muslim woman: Individual Religiosity: 0.1, Vertical Religious Pressure: 0.1, Local share of Maghrebi immigrants: .10, Muslim worship seats per thousand inhabitants: 2, Schooling: 15 years, Work experience: 4 years. One can compare those values with the summary statistics of Tables 2.6 and 2.7 to verify that this roughly corresponds to an average Muslim woman in our sample.

[23]To compute this counterfactual, we shut down the economic discrimination channel, and set the share of local Maghrebi immigrants to 0.6 (instead of 0.1) and the number of worship seats to 4 (instead of 2).

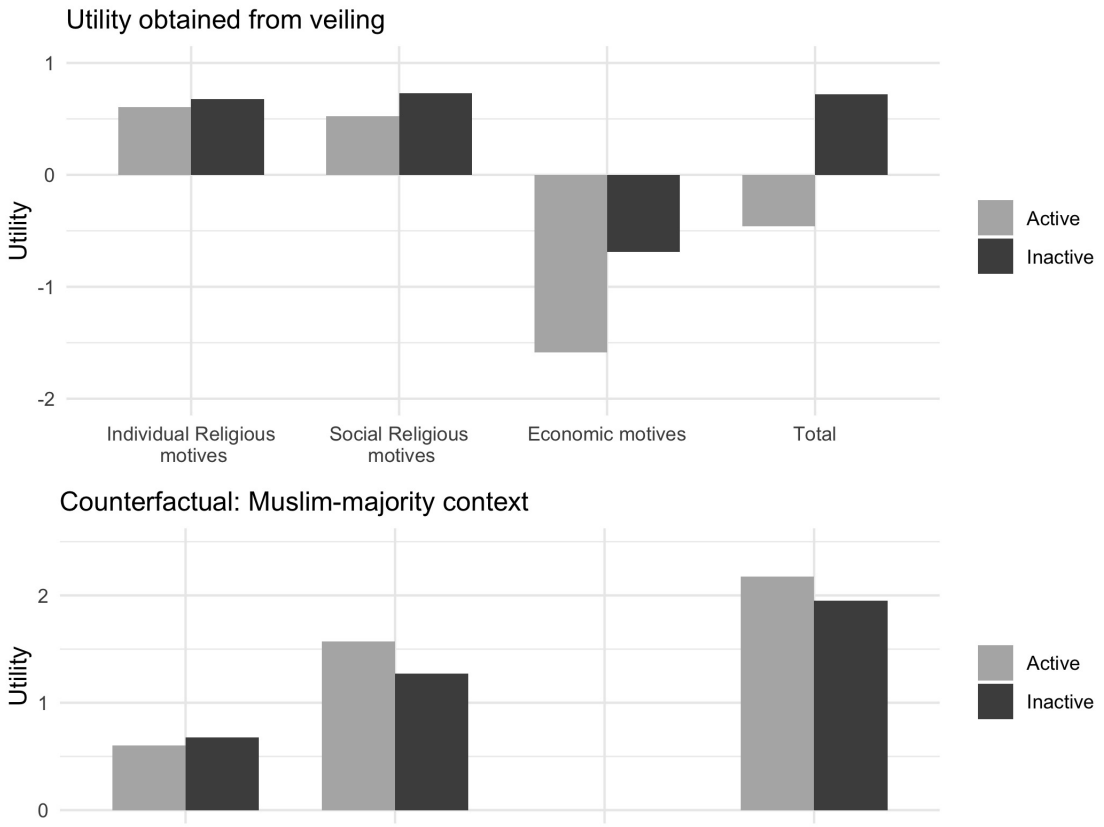

Figure 2.2: Utility obtained from veiling (i.e. difference of utility between choosing $v = 0$ and $v = 2$) according to the estimates of Table 2.5, and based on an 'average' woman in our sample (cf. footnote 22). The lower panel is obtained by counterfactual, shutting down the economic discrimination channel and modifying some characteristics of this average woman to reflect a Muslim-majority environment (cf. footnote 23).

(2013) model and with the evidence from Shofia (2020) on Indonesia, although the difference here is small.

In the rest of this section, we verify these results formally using the tests formulated in Implications 1–4. Detailed results for these tests are available in Appendix 2.7.1.

Baseline implications. Implications 1 and 2 concern the direction of variation for the coefficients β_{jv} and γ_{jv}, respectively with the veiling level v and the activity j. Tests of these implications should indicate whether our joint outcomes react to our predictors in the direction expected by the model.

Implication 1. Our first model implication concerns the relationship of our predictor variables with veiling behavior, within a given economic activity. Consider for instance our measure of individual religiosity. We can see clearly from Table 2.5 that individual religiosity is associated

with an increase in the degree of veiling, both for active and inactive women. Indeed, at activity j fixed, our estimates for β^1_{jv} increase across veiling levels v, indicating that higher measures of individual religiosity are associated with an increased propensity to wear the veil.

To verify this formally, we conduct hypothesis tests of the form $\beta^1_{jv'} - \beta^1_{jv} > 0$ for the different possible combinations of j, v and v' such that $v' > v$. (We present the detailed results in Figure 2.1, Appendix 2.7.1.) In this case, we find that Implication 1 holds at the 95% confidence level for all possible combinations of v and v', thus confirming the positive association between individual religiosity and veiling.

We then perform similar tests of Implication 1 for our five other main predictors. Most of our point estimates for the tests associated with the different predictors agree with Implication 1, although several tests do not reach statistical significance. Regarding vertical social pressure, five estimates out of six fall in the predicted region. For our first measure of horizontal pressure, i.e. the percentage of people from Maghrebi origin in the neighborhood, again five out of six point estimates fall in the predicted region. For our second measure of horizontal pressure, i.e. the local number of seats in religious facilities per 1000 inhabitants, four out of six point estimates fall in the predicted region, with two of those being significant at the 95% confidence level. Finally, both for our work experience variable and for our schooling variable, five out of six point estimates fall in the predicted region, with three of those being significantly different from zero.

Furthermore, if we ignore the 'discreet symbols' veiling category for which we have few observations, then our point estimates systematically fall in the half-space predicted by the model, with a majority of the tests yielding statistically significant predictions.

Put together, we interpret these results as providing partial evidence for Implication 1. Although a majority of the tests do not hold at the 95% level, the overall pattern of point estimates falling in the predicted region suggests some validity for the statement of Implication 1. Notably, statistical power might be an issue here, as we observed by discarding the estimates linked to the 'discreet symbols' category, for which we have few observations: doing so decreases the rejection rate for our tests. Overall, the tests of Implication 1 thus confirm that our religiosity variables are broadly associated with an increased propensity to veil, while our economic variables are associated with a decreased propensity to veil.

Implication 2. Our second implication concerns the relationship of our predictor variables with economic activity, holding the degree of veiling fixed. As we did with Implication 1, we perform tests of Implication 2 for our six main predictors, the results of which are presented in Figure 2.2. First, regarding our four religiosity variables, there does not seem to be much support for

Implication 2. There is no systematic pattern for point estimates as we observed for Implication 1, and all tests fail at the 95% confidence level. Therefore, we do not find any evidence for our religious variables being associated with an increased or decreased propensity to be economically active.

On the contrary, we find that our economic variables are strongly associated with economic activity. Indeed, Implication 2 holds for both our work experience and schooling variables. This indicates a strong positive association between these economic variables and the propensity to be economically active.

Since we do not find that religiosity variables are strongly associated with the propensity to be economically active, the 'religious motives channel' is already undermined by the tests of Implication 2. This is because this channel predicts that, when holding the degree of veiling constant, women who are more religious or who face more external religious pressure should be less economically active. However, this is not what we find here: our results suggest that the religiosity variables do not have a direct effect on economic participation, but only an indirect one through the practice of veiling. We discuss this further with the test of Implication 3 below.

Mechanisms. We now move on to the tests of Implications 3 and 4, which are more directly related to the two mechanisms that we highlighted above: the religious motives channel, and the economic discrimination channel.

Implication 3. Our third implication can be interpreted as a formal test for the religious motives channel, since it examines whether veiling has higher religious returns for women who are economically active, compared to those who are not. Our results for these tests are presented in Figure 2.3. In this case, neither test significance nor point estimates suggest that the formal statement of Implication 3 holds. As such, we do not find evidence for this mechanism.

This result is in line with those of the tests for Implication 2, which already suggested an absence of association between our religious variables and economic participation among Muslim women. Taken together, these results point towards religious motives having an effect on economic participation only through the practice of veiling. This supports the idea that the negative correlation between veiling and economic participation that we observed in the descriptive analysis may be mostly due to veiling having a cost on the labor market, as opposed to religious women having different preferences from non-religious women regarding economic participation.

Implication 4. Finally, our fourth implication can be interpreted as a formal test for the economic discrimination channel, by examining whether economic participation has higher returns for women who do not veil, compared to those who do. Results are presented in Figure 2.4.

Regarding our first economic variable, work experience, we do not find support for the statement of Implication 4: the tests reject the hypothesis at the 95% confidence level, and there is no pattern of point estimates mostly belonging to the predicted region. This is perhaps because, on average, work experience does not substantially differs by veiling status (see Table 2.6) since veiled women are older and thus had more time to accumulate experience. However, we find some support in the tests associated with our second economic variable, schooling, which most women in our sample had time to complete. In this case, all point estimates fall within the predicted region. Furthermore, the test which ignores the 'discreet symbols' category suggest statistically significant differences (although those which involve these categories do not hold at this level).

This second result offers support the economic discrimination channel: higher-educated women are less likely to integrate economically if they veil, even if we hold religiosity variables constant. In other words, the utility returns on schooling are lower for women who veil compared to those who do not. We have seen in our discussion of Implication 2 that this seems to be unrelated to an underlying preference towards economic participation linked with individual religiosity or social religious pressures. Therefore, this result seems to support the idea that there is an economic cost to veiling, in the sense that veiled women face weaker economic opportunities than those who do not veil.

To sum up, our results suggest that the interaction between the decision to veil and that of economic participation is mostly driven by economic concerns. First, both religious motives and economic ones play important roles in the decision to veil. Second, while economic motives are strong drivers of economic participation, the same is not true for religious motives, suggesting that the veil itself (and not underlying religious preferences) is linked to decreased economic participation. Third, non-veiled women seem to enjoy higher economic returns on their education compared to veiled women (holding individual religiosity and social religious pressures fixed), as evidenced by their higher propensity to be economically active.

Overall, those results suggest that the religious mechanism suggested by Carvalho (2013) cannot fully explain veiling and economic participation patterns in France. Instead, the interaction between veiling and the economic incentives to economic participation, such as the discrimination against veiled women on the labor market, seems to play an important role in

this context. Furthermore, and of particular importance for the French debate, we note that individual religious motives turn out to be at least as important as communitarian influences in the decision to veil.

2.6 Conclusion

Theoretical and empirical studies of veiling in economics have so far mainly focused on Muslim-majority countries, perhaps because of the paucity of data on veiling in developed countries. With the rising immigration flows of Muslims to secular countries, getting a better understanding of why women veil is nonetheless crucial as many countries, of which France is maybe the most emblematic, limit the expression of religious faith in public.

In this paper, we tackle this question using rare rich observational data on Muslim women in France. The richness of the data notably allows us to distinguish between private and communitarian incentives to veil. We first document that in France, wearing conspicuous religious symbols is associated with a much lower economic integration for Muslim women. The magnitude of this relationship is large, comparable to having a child less than 4 years old for instance. Second, we find that, among the main incentives for veiling highlighted in the economic literature, the wearing of conspicuous symbols appears to be strongly driven by private religious motivations. Third, we find that the joint decision to veil and being economically active can be mostly explained by economic (dis)incentives. Our results thus suggest that the veiling mechanism proposed by Carvalho (2013) and evidenced in the context of Indonesia by Shofia (2020) may be second-order in a non-Muslim-majority country such as France. Instead, when choosing whether to work and to wear the veil, Muslim women seem to be more sensitive to incentives related to how veiling impacts their economic opportunities.

Because they underline the role of private religious motives instead of community pressure ones, our results question the rhetoric often used to justify policies restricting the wearing of religious symbols in France. In the media and in political spheres, journalists and politicians almost always defend veiling restrictions on the basis that Muslim women are being forced to veil by their husband and community. If these claims were true, it is believed that secular policies could have the potential to "free" Muslim women from religious pressures and promote gender equality (e.g. Maurin and Navarrete-Hernandez 2023). Actually, even in this case, Carvalho (2013) shows that banning the wearing of the veil in public might lead to *more* segregation because women would lose the ability to signal their piety to their community. However, consistent with existing evidence from qualitative interviews with Muslim women, we find that the main incentives for veiling appear to be private. In other words, Muslim women who veil do so for

personal reasons linked to their own beliefs, first and foremost. Therefore, further restricting the wearing of conspicuous religious symbols is likely to lead to even poorer integration of Muslim women if these private benefits are high and discreet symbols are imperfect substitutes. Our complementary analysis of the Turkish case, a country which also imposed secular constraints in the public sphere, is consistent with this argument.

Furthermore, our results call attention to the importance of the discriminations that women who wear the veil face on the labor market. For instance, hiring discriminations against people who signal their Muslim affiliation were already documented by Valfort (2020). Because we find that individual religiosity and other religious factors seem to be associated with the decision to be economically active mainly through the act of veiling, a possible interpretation is that women who veil are less economically active not because of underlying preferences linked with their religiosity, but rather because the veil represents an obstacle to economic participation.

Our empirical approach in this paper is descriptive and should not be interpreted as causal. Still, our results suggest that veiling in France entails significant costs to economic integration, is driven by private incentives before social ones. Given the importance of better integrating Muslim populations in developed countries, future work could provide more robust assessments of the patterns uncovered in this paper. For example, if larger databases on Muslim women become available, one could evaluate the effect of external shocks to the local religious composition, such as exogenous migration waves, on veiling patterns. We finally note that data limitations inherent to studies of this type call for more initiatives like the TeO survey to better document the experiences of minority populations in a context of increasing global migrations.

Bibliography

Abdelgadir, A. and Fouka, V. (2020), 'Political secularism and Muslim integration in the West: Assessing the effects of the French Headscarf Ban', *American Political Science Review* **114**(3), 707–723.

Abdelhadi, E. (2019), 'The hijab and muslim women's employment in the united states', *Research in Social Stratification and Mobility* **61**, 26–37.

Abramitzky, R., Boustan, L. and Eriksson, K. (2020), 'Do immigrants assimilate more slowly today than in the past?', *American Economic Review: Insights* **2**(1), 125–41.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2010), 'Identifying barriers to muslim integration in france', *Proceedings of the National Academy of Sciences* **107**(52), 22384–22390.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2016), *Why Muslim integration fails in Christian-heritage societies*, Harvard University Press.

Aizer, A. and Cunha, F. (2012), The production of human capital: Endowments, investments and fertility. NBER Working Paper 18429.

Ajbli, F. (2011), Les Françaises musulmanes face à l'emploi: le cas des pratiquantes "voilées" dans la métropole lilloise, PhD thesis, Paris, EHESS.

Akerlof, G. A. and Kranton, R. E. (2000), 'Economics and identity', *The quarterly journal of economics* **115**(3), 715–753.

Aksoy, O. (2017), 'Motherhood, sex of the offspring, and religious signaling', *Sociological Science* **4**, 511–527.

Aksoy, O. and Gambetta, D. (2016), 'Behind the veil: the strategic use of religious garb', *European Sociological Review* **32**(6), 792–806.

Aksoy, O. and Gambetta, D. (2021), 'The politics behind the veil', *European Sociological Review* **37**(1), 67–88.

Algan, Y., Malgouyres, C., Mayer, T. and Thoenig, M. (2022), 'The economic incentives of cultural transmission: Spatial evidence from naming patterns across france', *The Economic Journal* **132**(642), 437–470.

Andriantsimbazovina, J., Kabou, P. et al. (2020), *Laïcité et défense de l'État de droit*, Presses de l'Université Toulouse 1 Capitole.

Atasoy, Y. (2006), 'Governing women's morality: a study of islamic veiling in Canada', *European Journal of Cultural Studies* **9**(2), 203–221.

Atkin, D., Colson-Sihra, E. and Shayo, M. (2021), 'How do we choose our identity? a revealed preference approach using food consumption', *Journal of Political Economy* **129**(4), 1193–1251.

Beauchemin, C., Hamel, C. and Simon, P. (2016), *Trajectoires et origines: enquête sur la diversité des populations en France*, INED éditions.

Becker, G. S. (1965), 'A theory of the allocation of time', *The Economic Journal* **75**(299), 493–517.

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

Becker, G. S. (1974), 'A theory of marriage: Part II', *Journal of Political Economy* **82**(2), S11–S26.

Becker, S. O., Grosfeld, I., Grosjean, P., Voigtländer, N. and Zhuravskaya, E. (2020), 'Forced migration and human capital: Evidence from post-WWII population transfers', *The American Economic Review* **110**(5), 1430–63.

Becker, S. O. and Woessmann, L. (2009), 'Was Weber wrong? A human capital theory of Protestant economic history', *The Quarterly Journal of Economics* **124**(2), 531–596.

Bellemare, M. F., Novak, L. and Steinmetz, T. L. (2015), 'All in the family: Explaining the persistence of female genital cutting in west africa', *Journal of Development Economics* **116**, 252–265.

Benjamini, Y. and Hochberg, Y. (1995), 'Controlling the false discovery rate: a practical and powerful approach to multiple testing', *Journal of the Royal statistical society: series B (Methodological)* **57**(1), 289–300.

Benzer, T. (2022), Removing cultural barriers to education: State-run islamic schools and girls' education in turkey, Working paper.

Berndt, E., Hall, B., Hall, R. and Hausman, J. (1974), 'Estimation and inference in nonlinear structural models', *Annals of Economic and Social Measurement* **3**(4), 653–665.

Bessin, M. (2002), *Autopsie du service militaire 1965–2001*, Autrement.

Biavaschi, C., Giulietti, C. and Siddique, Z. (2017), 'The economic payoff of name americanization', *Journal of Labor Economics* **35**(4), 1089–1116.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2011), 'Ethnic identity and labour market outcomes of immigrants in europe', *Economic Policy* **26**(65), 57–92.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2016), 'Bend it like beckham: Ethnic identity and integration', *European Economic Review* **90**, 146–164.

Bisin, A. and Topa, G. (2003), 'Empirical models of cultural transmission', *Journal of the European Economic Association* **1**(2/3), 363–375.

Bisin, A., Topa, G. and Verdier, T. (2004), 'Religious intermarriage and socialization in the United States', *Journal of Political Economy* **112**(3), 615–664.

Bisin, A. and Tura, G. (2022), Marriage, fertility, and cultural integration in Italy. Working Paper.

Bisin, A. and Verdier, T. (2000), '"Beyond the Melting Pot": Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits', *The Quarterly Journal of Economics* **115**(3), 955–988.

Bisin, A. and Verdier, T. (2001), 'The economics of cultural transmission and the dynamics of preferences', *Journal of Economic Theory* **97**, 298–319.

Bisin, A. and Verdier, T. (2011), The economics of cultural transmission and socialization, *in* 'Handbook of Social Economics', Vol. 1, North-Holland, pp. 339–416.

Bisin, A. and Verdier, T. (2022), Advances in the economic theory of cultural transmission. NBER Working Paper 30466.

Bolt, U., French, E., Maccuish, J. H. and O'Dea, C. (2021), The intergenerational elasticity of earnings: exploring the mechanisms, Cepr discussion paper no. dp15975.

Botticini, M. and Eckstein, Z. (2007), 'From farmers to merchants, conversion and diaspora: Human capital and Jewish history', *Journal of the European Economic Association* **5**(5), 885–926.

Botticini, M. and Eckstein, Z. (2012), *The Chosen Few*, Princeton University Press.

Bourdieu, P. (1979), 'Les trois états du capital culturel', *Actes de la recherche en sciences sociales* **30**(1), 3–6.

Bouzar, D. and Bouzar, L. (2009), *Allah a-t-il sa place dans l'entreprise?*, Albin Michel.

Carvalho, J.-P. (2013), 'Veiling', *The Quarterly Journal of Economics* **128**(1), 337–370.

Carvalho, J.-P., Koyama, M. and Sacks, M. (2017), 'Education, identity, and community: lessons from the Jewish emancipation', *Public Choice* **171**, 119–143.

Carvalho, J.-P., Koyama, M. and Williams, C. (2022), Resisting education. Oxford Discussion paper 982.

Carvalho, J.-P. and McBride, M. (2022), The formation of religious beliefs and preferences, *in* K. Zimmerman, ed., 'Handbook of Labor, Human Resources and Population Economics', Springer.

Chaudhary, L. and Rubin, J. (2011), 'Reading, writing, and religion: Institutions and human capital formation', *Journal of Comparative Economics* **39**, 17–33.

Cheung, M.-W. and Wu, J. (2018), 'On the probabilistic transmission of continuous cultural traits', *Journal of Economic Theory* **174**, 300–323.

Chiappori, P.-A. (1992), 'Collective labor supply and welfare', *Journal of Political Economy* **100**(3), 437–467.

Chiappori, P.-A. (2017), *Matching with transfers: The economics of love and marriage*, Princeton University Press.

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

Choo, E. and Siow, A. (2006), 'Who marries whom and why', *Journal of Political Economy* **11'**(1), 175–201.

Cintas, C., Gosse, B. and Vatteville, É. (2012), Quand l'identité religieuse devient une préoccupation du management des ressources humaines, *in* 'Management et religion: Décryptage d'un lien indéfectible', EMS Editions, pp. 83–98.

Corekcioglu, G. (2021), 'Unveiling the effects of a headscarf ban: Evidence from municipal jobs in turkey', *Journal of Comparative Economics* **49**(2), 382–404.

Cunha, F. and Heckman, J. (2007), 'The technology of skill formation', *The American Economic Review* **97**(2), 31–47.

Del Boca, D., Flinn, C. and Wiswall, M. (2014), 'Household choices and child development', *The Review of Economic Studies* **81**(1), 137–185.

Del Boca, D., Flinn, C. and Wiswall, M. (2016), 'Transfers to households with children and child development', *The Economic Journal* **126**(596), F136–F183.

Del Bono, E., Francesconi, M., Kelly, Y. and Sacker, A. (2016), 'Early maternal time investments and early child outcomes', *The Economic Journal* **126**, F96–F135.

Droogsma, R. A. (2007), 'Redefining hijab: American muslim women's standpoints on veiling', *Journal of Applied Communication Research* **35**(3), 294–319.

Drydakis, N. (2013), 'The effect of ethnic identity on the employment of immigrants', *Review of Economics of the Household* **11**(2), 285–308.

Fan, X. and Wu, L. (2022), The shaping of a gender norm: Marriage, labor, and foot-binding in historical china, Working paper.

Fernández-Reino, M., Di Stasio, V. and Veit, S. (2022), 'Discrimination unveiled: A field experiment on the barriers faced by muslim women in germany, the netherlands, and spain', *European Sociological Review* .

Fouka, V. (2020), 'Backlash: The unintended effects of language prohibition in U.S. schools after World War I', *The Review of Economic Studies* **87**(1), 204–239.

Fryer, R. G. and Levitt, S. D. (2004), 'The causes and consequences of distinctively black names', *The Quarterly Journal of Economics* **119**(3), 767–805.

Galichon, A., Kominers, S. D. and Weber, S. (2019), 'Costly concessions: An empirical framework for matching with imperfectly transferable utility', *Journal of Political Economy* **127**(6), 2875–2925.

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

Garre, F. G., Vermunt, J. K. and Croon, M. A. (2002), 'Likelihood-ratio tests for order-restricted log-linear models: A comparison of asymptotic and bootstrap methods', *Metodologa de las Ciencias del Comportamiento* **4**(1), 1–18.

Gaspard, F. and Khosrokhavar, F. (1995), *Le foulard et la République*, FeniXX.

Gibbons, R. and Waldman, M. (2004), 'Task-specific human capital', *The American Economic Review* **94**(2), 203–207.

Greene, W. H. (2008), *Econometric analysis*, 6th edn, Pearson Education, Inc.

Gresh, A. (2020), 'Islam: A one-sided debate'.
URL: *https://orientxxi.info/dossiers-et-series/islam-a-one-sided-debate,4281*

Guiso, L., Sapienza, P. and Zingales, L. (2006), 'Does culture affect economic outcomes?', *Journal of Economic Perspectives* **20**(2), 23–48.

Gulesci, S., Jindani, S., La Ferrara, E., Smerdon, D., Sulaiman, M. and Young, H. (2021), A stepping stone approach to understanding harmful norms, CEPR Discussion Paper No. DP15776.

Harrison, K. A. (2016), 'Hiding under the veil of" dress policy": Muslim women, hijab, and employment discrimination in the united states', *Geo. J. Gender & L.* **17**, 831.

Heckman, J., Pinto, R. and Savelyev, P. (2013), 'Understanding the mechanisms through which an influential early childhood program boosted adult outcomes', *American Economic Review* **103**(6), 2052–86.

Hu, Z. and Wang, Z. (2021), Nutrition, labor supply, and productivity: Evidence from ramadan in indonesia, Working paper.

Iannaccone, L. R. (1990), 'Religious practice: A human capital approach', *Journal for the Scientific Study of Religion* **29**(3), 297–314.

Iannaccone, L. R. (1992), 'Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives', *Journal of Political Economy* **100**(2), 271–291.

Iannaccone, L. R. (1998), 'Introduction to the economics of religion', *Journal of Economic Literature* **36**(3), 1465–1495.

Institut Français d'Opinion Publique [IFOP] (2019), Etude auprès de la population musulmane en france, 30 ans après l'affaire des foulards de creil, Technical report.

Institut Montaigne (2014–2021), Religion au travail: croire au dialogue. baromètre du fait religieux en entreprise 2014–2021, Technical report.

Institut Montaigne (2016), A french islam is possible, Technical report.

Iyer, S. (2016), 'The new economics of religion', *Journal of Economic Literature* **54**(2), 395–441.

Jia, R. and Persson, T. (2021), 'Choosing ethnicity: The interplay between individual and social motives', *Journal of the European Economic Association* **19**(2), 1203–1248.

Jouili, J. S. (2020), *Pious Practice and Secular Constraints*, Stanford University Press.

Katz, D., Baptista, J., Azen, S. and Pike, M. (1978), 'Obtaining confidence intervals for the risk ratio in cohort studies', *Biometrics* **34**(3), 469–474.

Khosrokhavar, F. (2004), L'islam des jeunes filles en france, *in* 'Le foulard islamique en questions', Éditions Amsterdam, pp. 89–94.

Kirkebøen, L., Leuven, E. and Mogstad, M. (2022), 'College as a marriage market', *NBER Working Paper* (28688).

La Boussole (2004), *Annuaire des mosquées de France*.

Lecomte, J.-P. (2001), Représentations et réalités des fonctions sociales du service militaire dans la société française (1868–2001). Thèse de doctorat en Science politique soutenue en 2001 à Paris, Institut d'études politiques.

Livny, A. (2020), *Trust and the Islamic advantage: religious-based movements in Turkey and the Muslim world*, Cambridge University Press.

Maillard, D. (2017), *Quand la religion s'invite dans l'entreprise*, Fayard.

Maurin, E. and Navarrete-Hernandez, N. (2023), 'Behind the veil: the effect of banning the islamic veil in schools', *Economic Policy* **38**(113), 63–98.

McFadden, D. (1973), Conditional logit analysis of qualitative choice behavior, *in* P. Zarembka, ed., 'Frontiers in Econometrics', Academic Press, New York, chapter 4, pp. 105–142.

Montgomery, J. D. (2010), 'Intergenerational cultural transmission as an evolutionary game', *American Economic Journal: Microeconomics* **2**(4), 115–136.

Morsy, M. (1989), *Les femmes du Prophète*, FeniXX.

Munshi, K. (2011), 'Strength in numbers: Networks as a solution to occupational traps', *The Review of Economic Studies* **78**(3), 1069–1101.

Munshi, K. (2019), 'Caste and the Indian economy', *Journal of Economic Literature* **57**(4), 781–834.

Nordmann, C. (2004), *Le foulard islamique en questions*, Éditions Amsterdam.

Novak, L. (2020), 'Persistent norms and tipping points: The case of female genital cutting', *Journal of Economic Behavior & Organization* **177**, 433–474.

Patacchini, E. and Zenou, Y. (2016), 'Social networks and parental behavior in the intergenerational transmission of religion', *Quantitative Economics* **7**, 969–995.

Patel, D. S. (2012), 'Concealing to reveal: The informational role of islamic dress', *Rationality and Society* **24**(3), 295–323.

Platteau, J.-P. (2017), *Islam instrumentalized*, Cambridge University Press.

Read, J. G. and Bartkowski, J. P. (2000), 'To veil or not to veil? a case study of identity negotiation among muslim women in austin, texas', *Gender & society* **14**(3), 395–417.

Roth, A. E. and Sotomayor, M. A. O. (1990), *Two-sided matching: A study in game-theoretic modeling and analysis*, Cambridge University Press.

Sakalli, S. O. (2019), Secularization and religious backlash: Evidence from turkey, Technical report, Technical Report, Working Paper.

Saleh, M. (2016), 'Public mass modern education, religion, and human capital in twentieth-century Egypt', *The Journal of Economic History* **76**(3), 697–735.

Saleh, M. (2018), 'On the road to heaven: Taxation, conversions, and the Coptic–Muslim socioeconomic gap in Medieval Egypt', *The Journal of Economic History* **78**(2), 394–434.

Scott, J. W. (2009), *The politics of the veil*, Vol. 7, Princeton University Press.

Sénat (2021), 'Compte rendu intégral: Séance du mardi 30 mars 2021', *Journal Officiel de la République française* .

Shofia, N. M. (2020), Why veil? religious headscarves and the public role of women, Working paper.

Silhouette-Dercourt, V., Sy, O. S. and Desjeux, D. (2019), 'Cosmopolitan veiling in paris: Young french muslim women in transition', *Youth and Globalization* **1**(1), 65–87.

Siow, A. (2015), 'Testing Becker's theory of positive assortative matching', *Journal of Labor Economics* **33**(2), 409–441.

Skrondal, A. and Rabe-Hesketh, S. (2009), 'Prediction in multilevel generalized linear models', *Journal of the Royal Statistical Society Series A: Statistics in Society* **172**(3), 659–687.

Squicciarini, M. (2020), 'Devotion and development: Religiosity, education, and economic progress in 19th-century France', *The American Economic Review* **110**(11), 3454–3491.

Valfort, M.-A. (2020), 'Anti-muslim discrimination in france: Evidence from a field experiment', *World Development* **135**, 105022.

van der Hasselt, G. (2019), 'The muslim veil in france: Why so controversial? three questions to hakim el kaouri'.
URL: *https://www.institutmontaigne.org/en/blog/muslim-veil-france-why-so-controversial*

Weber, M. (2013), *The Protestant Ethic and the Spirit of Capitalism*, Routledge.

Weichselbaumer, D. (2020), 'Multiple discrimination against female immigrants wearing headscarves', *ILR Review* **73**(3), 600–627.

2.7 Appendix

2.7.1 Data and additional results

Measurement of individual religiosity and communitarian pressures

The TeO dataset contains rich information on respondents' religious life. We first describe the variables we use to proxy for individual religiosity, vertical religious influence (from parents), and horizontal pressures (from Muslim peers). We then detail how we combine those multiple measures into meaningful indices through a measurement system.

Individual religiosity. In TeO1, we measure individual religiosity using survey questions on the frequency of attendance of religious ceremonies, the self-reported importance of religion in the respondent's life, whether she uses her religion to self-identify, the respect of religious dietary restrictions, and religious marriage. In TeO2, an additional variable is available, that is, the frequency of praying. We list details of these variables below:

Variable name	Values	Question	Type
attendance of religious ceremonies	never; for familial ceremonies only; for religious feasts only; one or twice a month; weekly	"How often do you attend religious ceremonies?"	ordinal
importance of religion in respondent's life	no importance; a little; quite important; very important	"What importance do you give to religion in your life today?"	ordinal
uses religion to self-identify	yes; no	"Among the following characteristics, which ones define you best? [...] Your religion?"	indicator
respect of dietary restrictions	never; sometimes; always; none (coded as a dummy if "always")	"In your daily life, do you respect your religion's dietary restrictions?"	indicator
religious marriage	yes; no	"Did you and your husband do a religious wedding?"	indicator

Vertical religious pressure. We measure vertical religious pressures using two variables, namely the self-reported importance of religion in the respondent's education and religious name-giving.

Horizontal religious pressure. We measure horizontal religious pressures (from Muslim peers) using two variables, namely the share of Maghrebi immigrants in the respondent's neighborhood (IRIS) and the local capacity in Muslim places of worship. In TeO1, the share of

Variable name	Values	Question	Type
importance of religion in education	no importance; a little important; quite important; very important	"What importance did religion have in the education you received in your family?"	ordinal
religious first name	yes; no	constructed by authors using respondent's first name	indicator

Maghrebi immigrants is reported in deciles of the distribution across France. We select the middle point of each bin, except for the extremes – zero or above 40%, where we set the value of the variable to 0 and 0.4 respectively. Our second proxy of local Muslim presence is the estimated capacity (by the Muslim association who produced the inventory) in Muslim places of worship at the local level. In TeO1, this is measured at the *commune* (municipal) level of residence for all French cities except Paris, Lyon, and Marseille, for which we observe the *arrondissement*.

Measurement system. For the first two concepts above, since there is no natural way to combine the ordinal and indicator variables into meaningful indices, we formulate a measurement system. We are interested in two latent variables, *individual religiosity* and *vertical religious pressure*, which we assume load into their respective proxies listed above. We interpret those proxies as noisy measures of the associated unobserved, underlying concept. Denote by Z and W the vectors of proxies for individual religiosity and for vertical pressure respectively. We assume ordinal relationships between measures $\{Z, W\}$ and underlying factors $\text{IndivReligiosity}_i$ and VertPressure_i:

$$Z_{i,j} = \mu^z_{1,j} + \lambda^z_j \, \text{IndivReligiosity}_i + \varepsilon^z_{i,j} \tag{2.13}$$

$$W_{i,j} = \mu^w_j + \lambda^w_j \, \text{VertPressure}_i + \varepsilon^w_{i,j} \tag{2.14}$$

where ε are measurement errors assumed to be i.i.d. and to follow a logistic distribution. As the latent factors do not have a natural scale or location, to simplify interpretations, we normalize the means of $\text{IndivReligiosity}_i$ and VertPressure_i to zero, and their variances to one. We then predict the latent factors for each individual by calculating their empirical Bayes means (Skrondal and Rabe-Hesketh 2009).

Summary statistics (TeO)

We present some novel summary statistics of Muslim women by veiling status in Table 2.6. We distinguish between four categories for the wearing of religious symbols, which depend on (1) whether the symbol is "discreet" or "conspicuous", and (2) whether it is worn "sometimes" or "always". Since there is very little variation in the number of symbols worn (most women report

only wearing one), we do not use that information and focus on the extensive margin. Along with the outside option of not wearing any symbol, we thus compare five veiling levels. In terms of the theoretical model, we interpret the veiling level (v) as being increasing in the following order: no symbol ($v = 0$), sometimes worn, and always worn. Overall, Muslim women wearing conspicuous religious symbols differ from other Muslim women in many respects. For example, they are on average older, have more children, and are more likely to live in a couple. Moreover, while most Muslim women wearing a discreet symbol are second-generation immigrants, the vast majority of women who wear a conspicuous symbol are first-generation immigrants. In line with a potential learning of the French social norms by women wearing discreet signs compared to those wearing the veil, the former are more likely to report being discriminated against for non-religious reasons, not to trust the French institutions, and to believe that racism is widespread in France.

In Table 2.6, we report summary statistics of all religion-related variables by veiling status. As expected, as we move toward “higher" veiling status, individuals report higher degrees of religiosity and live in more religious environments. For example, 79% of women who always wear conspicuous symbols report that religion is very important in their life, while less than half of women not wearing a religious symbol do so. Women wearing discreet symbols appear to be moderately religious, but still report higher degrees of religiosity than women without any symbol. Women who wear conspicuous symbols also seem to live in more religious environments: they are more likely to have a Muslim partner and to report that most of their friends are Muslims. Moreover, they live in communes (and neighborhoods) populated by a larger Muslim community (proxied by Maghrebi immigrants and Muslim places of worship). Veiled women also seem to be subject to stronger parental religious pressures. They are significantly more likely to report that religion was very important in their education and to be given a religious first name. In short, all of the core potential mechanisms mentioned so far display some association with veiling behavior in the expected directions (see Table 2.4).

The main fact that motivates the first part of our analysis is that women wearing religious symbols, in particular those who always do so, have much poorer labor-market and schooling outcomes than the rest of the sample. Indeed, women who always wear conspicuous religious symbols are much less economically active on average. Our measure of economic activity is the activity rate, that is, whether the woman is either working, studying, or looking for a job (unemployed) at the time of the survey. While less than 20% of women not wearing conspicuous signs are inactive at the time of the interview, this proportion increases to 30% for women who sometimes wear a conspicuous symbol and up to 64% for women who always do. Moreover,

while 20% of women not wearing a symbol report having never worked in their life, almost half of women who always veil indicate having never entered the labor force. In terms of schooling outcomes, Muslim women who wear a conspicuous symbol are less likely to have any schooling degree. They have completed, on average, 2 to 7 fewer years of schooling than Muslim women who wear discreet symbols or none. Overall, the data suggests that wearing the veil seems to be strongly associated with a decline in economic integration, but this correlation may be due to many other factors over which veiled women differ from other Muslim women. We therefore provide a more thorough regression analysis of this pattern in our empirical approach.

Analysis of panel data

Exploiting the respondents' employment history available in the TeO data, we construct a retrospective panel dataset of economic activity to test the robustness of our results to the timing of the survey. We restrict the sample to adults, meaning that we remove observations for which an individual is aged less than 18 years old. This sample selection is made because it can be plausibly assumed that the veiling decision, on average, is made before adulthood.[24] We estimate random effects models using this data and report results in Table 2.8. In column (1), we regress the activity rate on veiling status and year fixed effects. In columns (2) and (3), we include, in turn, time-varying observables and time-invariant controls. The time-invariant controls are all covariates and dummies included in the cross-sectional analysis that are not likely to have changed over time (at least after age 18). These include the mother's and father's religion (Muslim or other), whether the individual has an Arabic-sounding name, attendance of religious ceremonies (proxy for religiosity), self-reported feelings of French identity, the importance of religion in the respondent's education, birthplace dummies, and a set of survey fixed effects. In these regressions, we cluster standard errors at the individual level to account for serial correlation. However, we cannot include individual fixed effects because we do not have panel data on veiling. We thus implicitly assume that the veiling decision is permanent, which we argue is a reasonable assumption because "unveiling" is a relatively rare phenomenon in France.[25]

[24] In the case of the Islamic veil, ethnographic evidence shows that the decision is usually made between the age of reaching puberty and around 20 years old (Gaspard and Khosrokhavar 1995). According to Islamic prescriptions, girls are supposed to dress modestly (including covering their hair) when reaching puberty so as to reduce men's temptation. In reality, in France, many adolescents or young women choose to veil a few years after reaching puberty, that is, around adulthood. We also verify that our results are not sensitive to the 18 years old threshold. In a robustness check, we restrict the sample to individuals aged at least 25 years old and find similar results.

[25] Two surveys conducted over (rather small) representative samples of the French Muslim population suggest that between 8 and 10 percent of women of Muslim faith declare having worn the veil in the past and are no longer doing so (IFOP 2019, Institut Montaigne 2016). Out of the total number of women not currently wearing the veil, this figure represents between 12.3% and 14.7%. Since here, we have both untreated individuals to which we assign treatment and treated individuals whom we assign to the

Table 2.6: Summary statistics by veiling status, Muslim women

Veiling status:	No symbol	Sometimes discreet	Always discreet	Sometimes consp.	Always consp.	Diff (C-D)
Demographics						
Age in 2008	35.55	28.40	25.06	35.94	36.00	8.62***
First-gen. immigrant	0.61	0.24	0.51	0.68	0.78	
Second-gen. immigrant	0.39	0.66	0.49	0.32	0.22	-0.46***
Number of children	1.78	1.11	0.63	2.26	2.79	1.88***
Lives in a couple	0.59	0.49	0.48	0.68	0.74	0.34***
Not a French speaker	0.07	0.02	0.01	0.14	0.32	0.26***
Labour-force status in 2008						
Employed	0.54	0.43	0.36	0.44	0.22	-0.17***
Unemployed	0.18	0.23	0.27	0.12	0.09	-0.10***
Inactive	0.19	0.15	0.24	0.30	0.65	0.44***
Student	0.09	0.20	0.13	0.14	0.03	-0.15***
Has never worked	0.19	0.29	0.48	0.31	0.50	0.16***
Schooling attainment and work experience						
Completed high school	0.78	0.85	0.58	0.68	0.61	-0.22***
Higher education degree	0.22	0.24	0.10	0.20	0.19	-0.06**
Years of schooling	15.30	17.41	15.69	12.86	11.11	-6.09***
Years of work experience	7.06	3.93	3.44	5.75	2.66	-0.61*
Social life and integration						
Participates in household's food shopping	0.49	0.39	0.34	0.59	0.69	0.30***
Often meets her family	0.89	0.89	0.89	0.89	0.93	0.03
Often meets her friends	0.88	0.90	0.94	0.87	0.90	-0.03
Meets with neighbors	0.41	0.45	0.50	0.52	0.62	0.13***
Meets with work colleagues[1]	0.32	0.36	0.33	0.22	0.11	-0.11**
Visits some recreation sites	0.67	0.78	0.76	0.53	0.42	-0.32***
Refuses to visit some recreation sites	0.09	0.12	0.15	0.06	0.04	-0.08***
Belongs to an association	0.17	0.18	0.21	0.18	0.12	-0.07**
Brings the children to school most of the time[1]	0.78	0.88	0.78	0.83	0.82	-0.02
Opinions on discrimination and French institutions						
Victim of racism due to religion	0.36	0.50	0.56	0.51	0.66	0.09***
Victim of racism due to origins	0.79	0.84	0.84	0.83	0.75	-0.07**
Victim of discrimination in past 5 years	0.28	0.41	0.34	0.40	0.28	-0.07**
Believes that racism happens often in France	0.49	0.60	0.68	0.45	0.38	-0.25***
Does not trust the French justice system	0.23	0.28	0.32	0.20	0.20	-0.10***
Does not trust the French police	0.29	0.40	0.50	0.28	0.25	-0.19***
Does not trust the French school	0.07	0.10	0.15	0.07	0.06	-0.06***
ID controlled by the police at least once	0.18	0.28	0.31	0.28	0.12	-0.14***
Observations	2,017	166	151	148	516	

Note: The data source is the Trajectories and Origins (TeO) dataset of 2008. Veiling status is measured using the respondents' answers to the wearing of religious symbols. We distinguish four categories depending on (1) whether the symbol is "discreet" or "conspicuous", and (2) whether it is worn "sometimes" or "always". In the last column, we report differences in means between individuals wearing conspicuous and those wearing discreet symbols where we pooled individuals along the first dimension (salience) as well as significance levels of those differences. Level of significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

[1] Meeting with work colleagues is conditional on employment and bringing children to school is conditional on having children. Thus, these variables are measured over restricted samples.

Table 2.7: Religious environment and religiosity by veiling status, Muslim women

Veiling status:	No symbol	Sometimes discreet	Always discreet	Sometimes consp.	Always consp.	Diff (C-D)
Religious environment						
Muslim partner	0.56	0.49	0.53	0.74	0.76	0.33***
Muslim father	0.94	0.95	0.68	0.96	0.98	0.05***
Muslim mother	0.94	0.95	0.75	0.99	0.97	0.06***
At least half of friends are Muslims	0.719	0.783	0.675	0.838	0.919	0.17***
At least half of work colleagues are immigrants[1]	0.43	0.37	0.42	0.46	0.55	0.14**
Had conflicts on religion with parents when 18 years old	0.15	0.17	0.19	0.07	0.11	-0.04*
Individual religiosity						
Importance of religion in one's life						
A little important	0.18	0.11	0.05	0.04	0.04	-0.06***
Quite important	0.31	0.30	0.29	0.33	0.15	-0.14***
Very important	0.47	0.58	0.64	0.63	0.81	0.20***
Attends religious ceremonies						
Familial ceremonies only	0.290	0.329	0.247	0.284	0.198	-0.07**
Religious feasts only	0.216	0.348	0.273	0.372	0.283	-0.01
Once or twice a month	0.036	0.061	0.047	0.088	0.099	0.05***
At least once a week	0.027	0.006	0.047	0.088	0.155	0.11***
Other indicators of religiosity						
Always respects the religious dietary restrictions	0.826	0.898	0.901	0.946	0.975	0.07***
Religious marriage	0.390	0.307	0.298	0.527	0.657	0.33***
Share of children with a religious first name[1]	0.030	0.013	0.096	0.172	0.186	0.06***
Uses her religion to self-identify	0.13	0.21	0.12	0.25	0.22	0.05*
Parental influence and communitarian religious presence						
Religious first name	0.09	0.08	0.04	0.18	0.13	0.05***
Local Front National vote share	0.098	0.100	0.099	0.102	0.106	0.005***
Importance of religion in education received						
A little important	0.173	0.115	0.139	0.068	0.074	-0.06***
Quite important	0.303	0.265	0.231	0.225	0.198	-0.05
Very important	0.468	0.566	0.543	0.674	0.708	0.14***
Percentage of Maghrebi immigrants in IRIS of residence						
(5.9%, 10.7%]	0.086	0.066	0.093	0.095	0.045	-0.02
(10.7%, 16.7%]	0.150	0.199	0.166	0.088	0.130	-0.06***
(16.7%, 27.3%]	0.289	0.295	0.265	0.304	0.275	0.00
More than 27.3%	0.418	0.398	0.417	0.473	0.510	0.09***
Presence of Muslim places of worship in commune (or arrond.)						
Places of worship (/1000 inh.)	0.053	0.047	0.050	0.055	0.069	0.01***
Capacity in a place of worship (/1000 inh.)	12.249	8.882	11.498	12.582	17.243	5.42***
Capacity for women in a place of worship (/1000 inh.)	2.061	1.600	2.197	2.041	3.095	0.94***
Observations	2,017	166	151	148	516	

Note: The data source is the Trajectories and Origins (TeO) dataset of 2008. Veiling status is measured using the respondents' answers to the wearing of religious symbols. We distinguish four categories depending on whether (1) the symbol is "discreet" or "conspicuous", and (2) it is worn "sometimes" or "always". In the last column, we report differences in means between individuals wearing conspicuous and those wearing discreet symbols where we pooled individuals along the first dimension (salience) as well as significance levels of those differences. Level of significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

[1] The composition of work colleagues is conditional on employment and names of the respondents' children is conditional on having children. Thus, these variables are measured over restricted samples.

The results from these regressions overall confirm the findings obtained in the cross-sectional analysis. Indeed, the wearing of a conspicuous symbol is associated with a significant decline in economic participation. Once more, the estimated effect is much stronger when the individual always wears the symbol. The estimates are smaller in magnitude then those obtained in the cross-section, but are still statistically and economically significant. The results indicate that women who always veil are 20 percentage points less likely to be active than women not wearing any religious symbol in a given year. Other important determinants of the activity rate, as expected, are the number of young children, marital status, and the number of years of schooling. These results suggest that those obtained in section 2.3.3 are not merely due to the timing of the survey and portray a more general phenomenon about Muslim women in France.

untreated group, it is not clear in which direction this measurement error biases our estimates. In light of those issues, we treat this analysis simply as a robustness check of our main results obtained in the cross-section.

Table 2.8: Effect of veiling on economic participation of adult Muslim women, retrospective panel data

Dep. variable: activity dummy	(1)	(2)	(3)	25 y.o. +
Veiling status				
Sometimes discrete	0.102***	0.002	0.006	-0.013
	(0.026)	(0.020)	(0.020)	(0.038)
Always discrete	0.077*	-0.031	-0.024	-0.050
	(0.030)	(0.021)	(0.021)	(0.039)
Sometimes conspicuous	-0.120***	-0.052*	-0.039	-0.046
	(0.035)	(0.026)	(0.026)	(0.036)
Always conspicuous	-0.365***	-0.216***	-0.176***	-0.203***
	(0.020)	(0.017)	(0.017)	(0.023)
Educational attainment				
Years of schooling in France		0.012***	0.010***	0.009***
		(0.001)	(0.001)	(0.001)
Years of schooling abroad		0.001	0.001	0.000
		(0.001)	(0.001)	(0.001)
Time-varying demographics				
Age		-0.010*	-0.008	0.020*
		(0.004)	(0.005)	(0.008)
Age squared		0.000	0.000	-0.000**
		(0.000)	(0.000)	(0.000)
Number of children		-0.007	-0.007	-0.022***
		(0.005)	(0.005)	(0.006)
Number of children below age 4		-0.089***	-0.089***	-0.066***
		(0.006)	(0.006)	(0.007)
Married		-0.147***	-0.139***	-0.068***
		(0.014)	(0.014)	(0.019)
Constant	0.629***	0.756***	0.928***	0.484*
	(0.019)	(0.074)	(0.108)	(0.234)
Time-invariant controls	N	N	Y	Y
Year fixed effects	Y	Y	Y	Y
Number of individuals	2,790	2,790	2,790	2,053
Total observations (N X Years)	37680	37680	37680	25354
R^2	0.124	0.394	0.405	0.345

This table shows the results of random-effects regression models of the economic activity dummy on the veiling status and other covariates in the retrospective panel dataset. Standard errors clustered at the individual level in parentheses. The estimation sample is restricted to adult Muslim women with no missing covariates and to time periods during which the individual was in France. In the last column, we estimate the specification in column (3) on the restricted sample of individuals aged at least 25 years old. Level of significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Multi-logit regressions with controls

In Table 2.9 we present results similar to those of Table 2.5, but including additional controls.

Table 2.9: Determinants of joint employment and veiling decision, multinomial logit.

Activity choice (j)	Inactive ($j=0$)			Active ($j=1$)		
Veiling choice (v)	None (*baseline*)	Discreet (1)	Conspicuous (2)	None (3)	Discreet (4)	Conspicuous (5)
Indiv. religiosity (β^1_{jv})	0	0.42* (0.24)	2.13*** (0.26)	0.19 (0.18)	1.06*** (0.22)	2.19*** (0.35)
Vert. pressure (β^2_{jv})	0	-0.39 (1.44)	1.84** (0.83)	0.61 (0.75)	1.61* (0.96)	1.66* (0.97)
Horiz. pressure						
ShareMaghrebi$_i$ (β^3_{jv})	0	3.59* (2.12)	0.85 (1.13)	0.01 (0.89)	0.08 (1.04)	2.35 (1.53)
CapacityMosques$_i$ (β^4_{jv})	0	-0.12* (0.07)	0.10*** (0.03)	0.01 (0.03)	-0.05* (0.03)	0.04 (0.03)
Schooling (γ^1_{jv})	0	-0.03 (0.03)	-0.05** (0.02)	0.07*** (0.02)	0.03 (0.03)	-0.02 (0.02)
Work experience (γ^2_{jv})	0	-0.09* (0.05)	-0.04 (0.03)	0.17*** (0.02)	0.17*** (0.03)	0.11*** (0.03)
Observations	2802					
Sampling weights	✓					
Additional controls[1]	✓					
Pseudo R^2	0.216					

Note: This table reports estimates of the parameters of the econometric model (2.11). The baseline category is the choice of inactivity and not wearing any religious symbol. Individual religiosity and vertical religious pressures are measured as indices (with mean zero and variance 1) constructed from multiple proxies available in the TeO data (see Appendix 2.7.1 for details). ShareMaghrebi$_i$ is the proportion of the local population that is of Maghrebi origin. CapacityMosques$_i$ is the estimated capacity in Muslim places of worship in the area of residence. Robust standard errors in parentheses. Level of statistical significance : * $p<0.1$, ** $p<0.05$, *** $p<0.01$.
[1] Additional controls include age, age squared, marital status (a dummy for having a partner), a dummy equal to one if the partner is working, immigration status and a set of dummy variables for quintiles of the local (neighborhood-level) unemployment rate of immigrants.

Plots for the tests of the four implications

In Figures 2.1 to 2.4 we present the results of the tests of Implications 1–4, respectively.

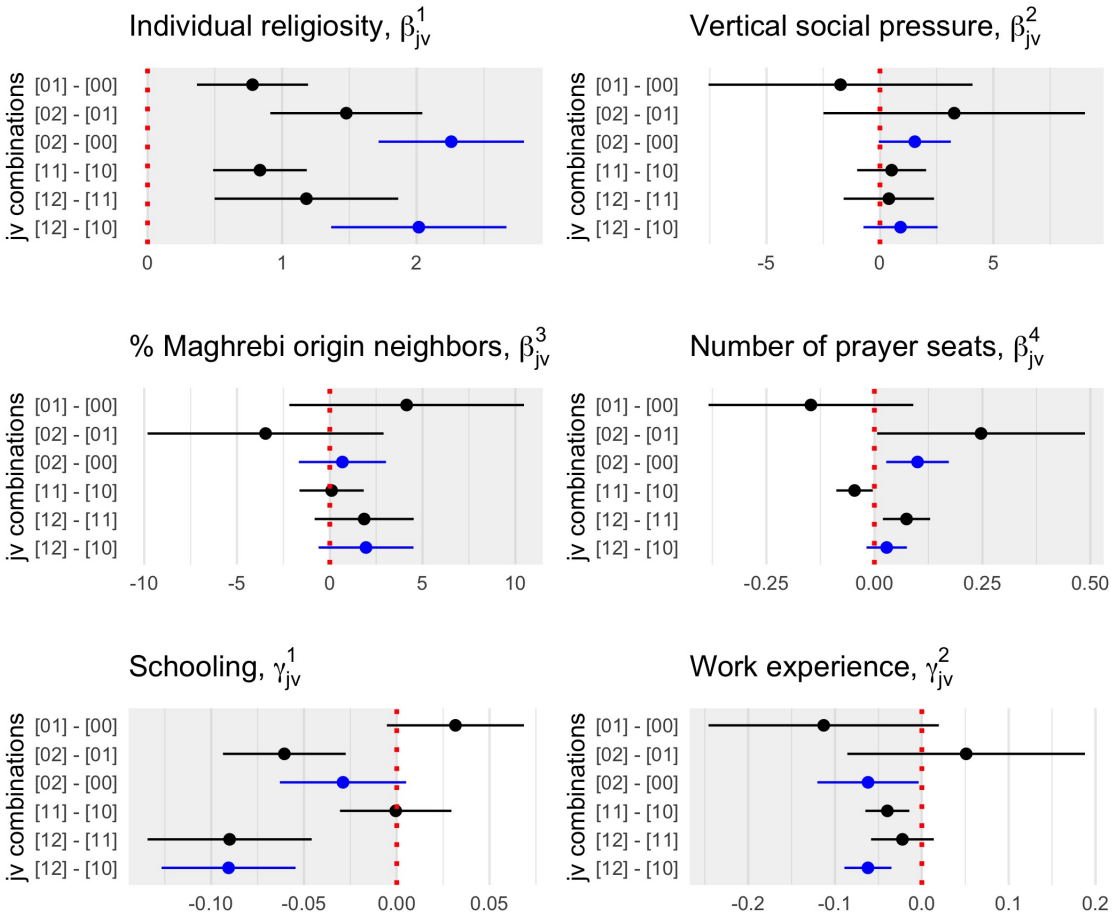

Figure 2.1: Hypothesis tests for Implication 1. Shaded areas correspond to the region where estimates are predicted to fall. Vertical axis labels correspond to the combination of (j, v) alternatives (e.g. the first line of the top-left graph plots the estimate for $\beta^1_{01} - \beta^1_{00}$). In blue: combinations which compare conspicuous symbol-wearing with no symbol-wearing. In black: combinations which include intermediate comparisons with discrete symbol-wearing. 95% confidence intervals are reported.

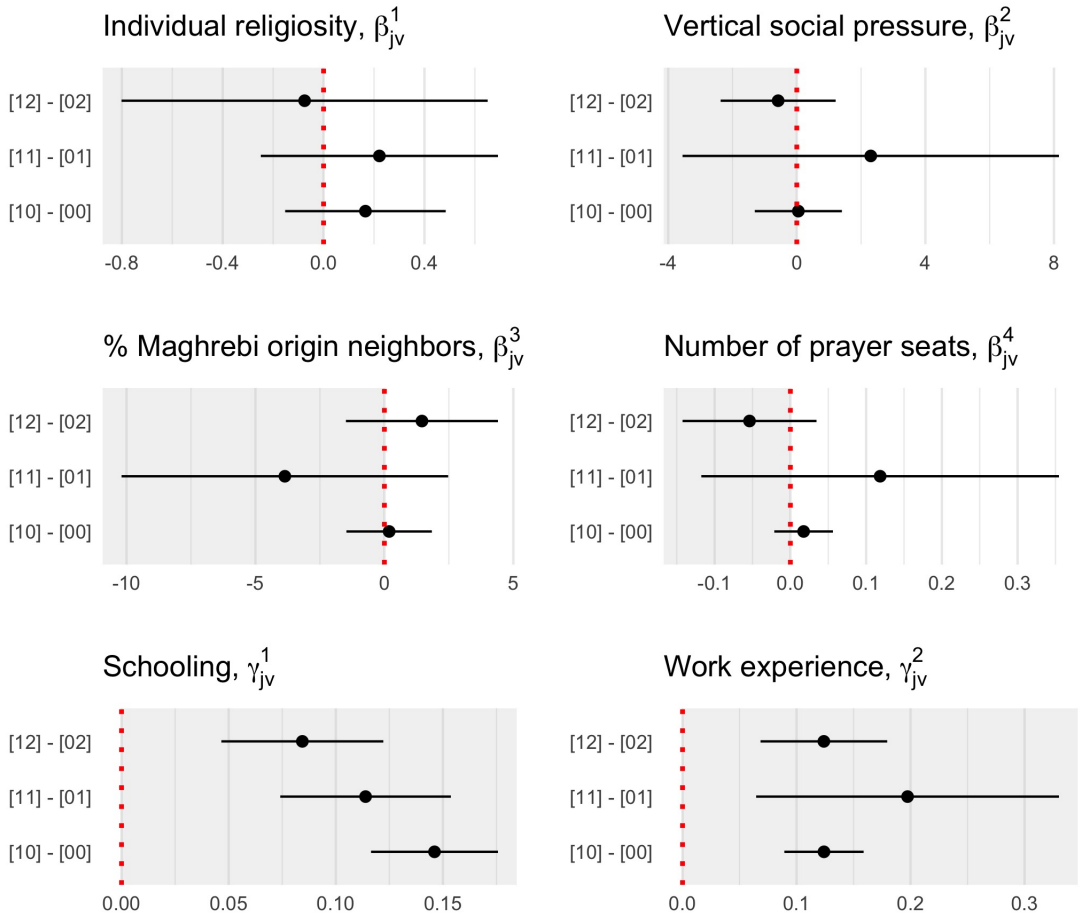

Figure 2.2: Hypothesis tests for Implication 2. Shaded areas correspond to the region where estimates are predicted to fall. Vertical axis labels correspond to the combination of (j, v) alternatives (e.g. the first line of the top-left graph plots the estimate for $\beta^1_{12} - \beta^1_{02}$). 95% confidence intervals are reported.

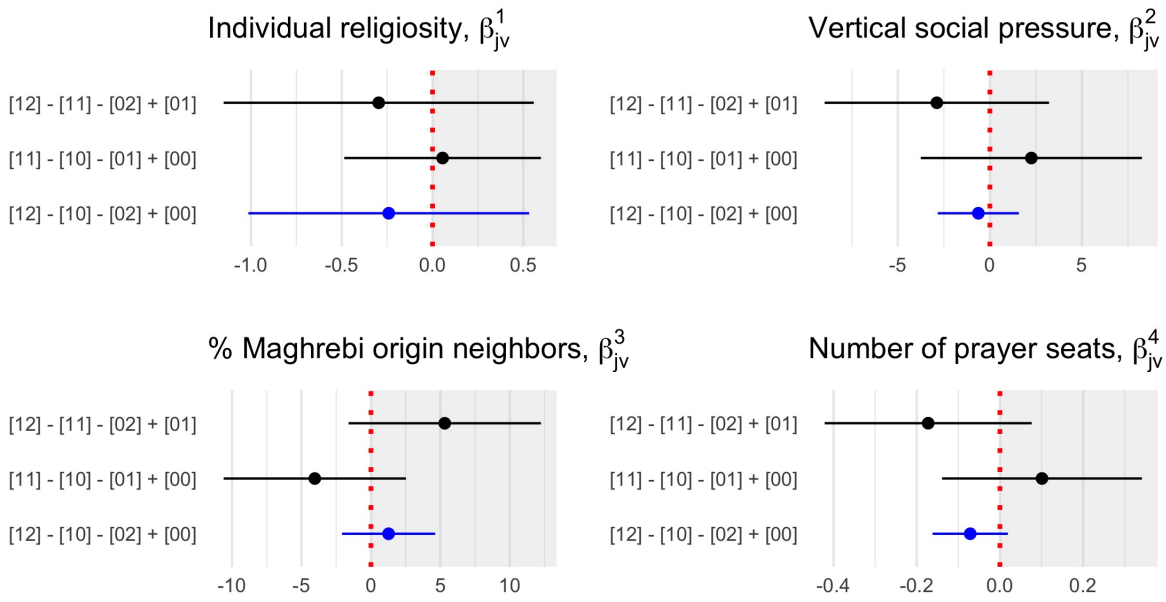

Figure 2.3: Hypothesis tests for Implication 3: Religious motives channel. Shaded areas correspond to the region where estimates are predicted to fall. Vertical axis labels correspond to the combination of (j, v) alternatives (e.g. the first line of the top-left graph plots the estimate for $\beta^1_{12} - \beta^1_{11} - \beta^1_{02} + \beta^1_{01}$). In blue: combinations which compare conspicuous symbol-wearing with no symbol-wearing. In black: combinations which include intermediate comparisons with discrete symbol-wearing. 95% confidence intervals are reported.

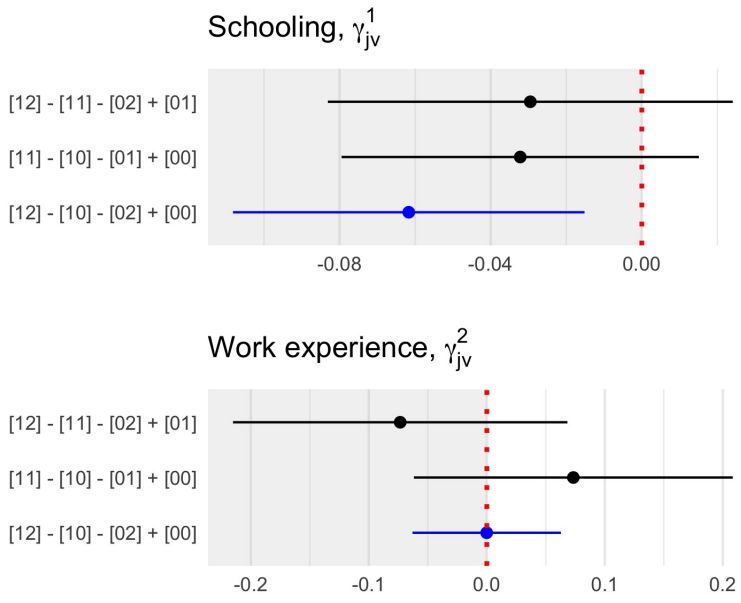

Figure 2.4: Hypothesis tests for Implication 4: Economic discrimination channel. Shaded areas correspond to the region where estimates are predicted to fall. Vertical axis labels correspond to the combination of (j, v) alternatives (e.g. the first line of the top-left graph plots the estimate for $\beta^1_{12} - \beta^1_{11} - \beta^1_{02} + \beta^1_{01}$). In blue: combinations which compare conspicuous symbol-wearing with no symbol-wearing. In black: combinations which include intermediate comparisons with discrete symbol-wearing. 95% confidence intervals are reported.

2.7.2 Veiling and economic outcomes in Turkey

In this Appendix, we explore the relationship between veiling and economic outcomes in Turkey and compare it to what we obtained for France and to that found by Shofia (2020) for Indonesia. Turkey is an interesting context to study veiling patterns since "it has long been considered a unique case of successful modernization through secularization" (Platteau 2017, p.355). Between the proclamation of the Turkish Republic, in October 1923, and the rise of the pro-Islamic conservative Justice and Development Party (AKP) to power in the early 2000s, the country was ruled by secular governments. The founders of the Republic implemented a top-down nationalist modernization project to "Westernize" Turkey. A major aspect of the multiple reforms adopted over the following decades was their secular nature as the government wanted to build a national identity that would subordinate the religious one (Sakalli 2019). Inspired by French State secularization, reforms ranging from the abolishment of the Caliphate to the adoption of Western dress codes profoundly changed the Turks' religious life. The series of secular legislation included veil bans in the public sphere. The 1982 Turkish constitution regulates veiling for civil servants, requiring women to uncover their head while on duty. The ban on headscarves was then extended to all universities in Turkey in 1997. Those regulations stayed in effect until they were gradually repealed by AKP: in 2010 for university campuses; in 2013 for state institutions; in 2014 for high schools; in 2016 for policewomen; and in 2017 for female army officers (Corekcioglu 2021).

Given that, despite the secular modernization of Turkey, Islam is by far the most prominent religion in the country, we see Turkey as an intermediate case between France and Indonesia in our theoretical framework. Similar to France, women face legal disincentives to veil in public. However, like Indonesia, Turkey is a Muslim-majority country. Therefore, we would expect the correlation between veiling and economic outcomes in Turkey to mirror those differences. Specifically, we expect the correlation between veiling and economic participation to be *negative*, but lower in magnitude than what we see in France because most of the Turkish society is religious.

To study the patterns of veiling and economic participation, we use Turkish data compiled from multiple sources by Livny (2020).[26] Importantly, these data contain information on veiling practices in Turkey, which is available at the district level. We collapse the different types of veils (turban, hijab, and burka) so as to obtain a single measure of veiling rate in each district. For economic outcomes, so as to harmonize those variables with our measures of veiling that span the years 2010 to 2015, we take the average of the outcomes in the district (province for GDP per capita) over the same time period. In Figure 2.1, we plot the relationship between the

[26]The data are publicly available on Avital Livny's website (https://www.alivny.com/data).

veiling rate and four measures of economic participation (female primary and secondary school completion, the female literacy rate, and GDP per capita) along with a quadratic fit.[27] For all of the outcomes we observe a negative association, suggesting that, in Turkey as in France, the veil might not act as an integration strategy. Interestingly, these negative relationships appear to be linear as most of the (small) curvature is driven by regions of the veiling-rate distribution with low mass (i.e. districts with low veiling rates).

We take these results as further suggestive evidence in line with the theory. The wearing of the veil was frowned upon by the secular elite before the bans were repealed, thus imposing a high cost to women when they veil and are economically active. Actually, as Platteau (2017) argues, the rise of an Islamist party to power reinforced the laicists' attachment to the secular values. Islamic identity signs, such as the veil, were sometimes also seen as manifesting a political identity in the public sphere in an increasingly polarized political context. Thus, even if Turkey is a Muslim-majority country, we find that the positive correlation documented by Shofia (2020) in Indonesia does not hold in this data. This suggests that her results regarding veiling behavior and economic participation are context-specific. Viewed through the lens of our theoretical framework, such a correlation can hold in Indonesia only because of two concomitant factors: (1) Indonesia is a Muslim-majority country, and (2) the veil is not subject to social or legal disapproval.

[27] For robustness, we also checked whether this relationship could be driven by religiosity of the district. We produced similar plots in which we control for religiosity and find very similar conclusions. Results are available upon request.

Figure 2.1: Relationship between veiling and economic outcomes at district level, Turkey 2010–2015

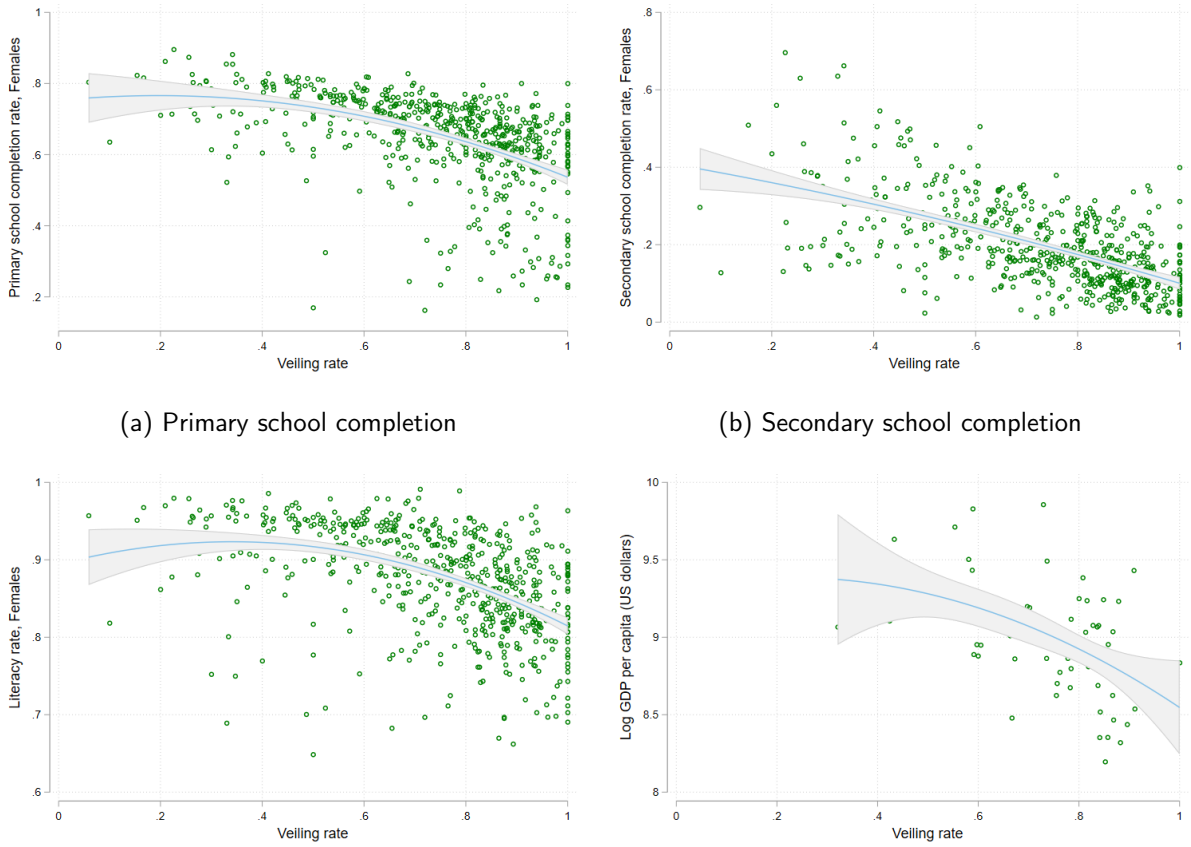

(a) Primary school completion

(b) Secondary school completion

(c) Literacy

(d) GDP per capita

Note: The data source is Livny (2020). These figures plot the relationship between the veiling rate in a district in 2010–2015 and the average of an economic outcome in that district over the same period, along with a quadratic fit and 95% confidence bands. For GDP per capita, the dependent variable is measured at the province level.

Bibliography

Abdelgadir, A. and Fouka, V. (2020), 'Political secularism and Muslim integration in the West: Assessing the effects of the French Headscarf Ban', *American Political Science Review* **114**(3), 707–723.

Abdelhadi, E. (2019), 'The hijab and muslim women's employment in the united states', *Research in Social Stratification and Mobility* **61**, 26–37.

Abramitzky, R., Boustan, L. and Eriksson, K. (2020), 'Do immigrants assimilate more slowly today than in the past?', *American Economic Review: Insights* **2**(1), 125–41.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2010), 'Identifying barriers to muslim integration in france', *Proceedings of the National Academy of Sciences* **107**(52), 22384–22390.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2016), *Why Muslim integration fails in Christian-heritage societies*, Harvard University Press.

Aizer, A. and Cunha, F. (2012), The production of human capital: Endowments, investments and fertility. NBER Working Paper 18429.

Ajbli, F. (2011), Les Françaises musulmanes face à l'emploi: le cas des pratiquantes "voilées" dans la métropole lilloise, PhD thesis, Paris, EHESS.

Akerlof, G. A. and Kranton, R. E. (2000), 'Economics and identity', *The quarterly journal of economics* **115**(3), 715–753.

Aksoy, O. (2017), 'Motherhood, sex of the offspring, and religious signaling', *Sociological Science* **4**, 511–527.

Aksoy, O. and Gambetta, D. (2016), 'Behind the veil: the strategic use of religious garb', *European Sociological Review* **32**(6), 792–806.

Aksoy, O. and Gambetta, D. (2021), 'The politics behind the veil', *European Sociological Review* **37**(1), 67–88.

Algan, Y., Malgouyres, C., Mayer, T. and Thoenig, M. (2022), 'The economic incentives of cultural transmission: Spatial evidence from naming patterns across france', *The Economic Journal* **132**(642), 437–470.

Andriantsimbazovina, J., Kabou, P. et al. (2020), *Laïcité et défense de l'État de droit*, Presses de l'Université Toulouse 1 Capitole.

Atasoy, Y. (2006), 'Governing women's morality: a study of islamic veiling in Canada', *European Journal of Cultural Studies* **9**(2), 203–221.

Atkin, D., Colson-Sihra, E. and Shayo, M. (2021), 'How do we choose our identity? a revealed preference approach using food consumption', *Journal of Political Economy* **129**(4), 1193–1251.

Beauchemin, C., Hamel, C. and Simon, P. (2016), *Trajectoires et origines: enquête sur la diversité des populations en France*, INED éditions.

Becker, G. S. (1965), 'A theory of the allocation of time', *The Economic Journal* **75**(299), 493–517.

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

Becker, G. S. (1974), 'A theory of marriage: Part II', *Journal of Political Economy* **82**(2), S11–S26.

Becker, S. O., Grosfeld, I., Grosjean, P., Voigtländer, N. and Zhuravskaya, E. (2020), 'Forced migration and human capital: Evidence from post-WWII population transfers', *The American Economic Review* **110**(5), 1430–63.

Becker, S. O. and Woessmann, L. (2009), 'Was Weber wrong? A human capital theory of Protestant economic history', *The Quarterly Journal of Economics* **124**(2), 531–596.

Bellemare, M. F., Novak, L. and Steinmetz, T. L. (2015), 'All in the family: Explaining the persistence of female genital cutting in west africa', *Journal of Development Economics* **116**, 252–265.

Benjamini, Y. and Hochberg, Y. (1995), 'Controlling the false discovery rate: a practical and powerful approach to multiple testing', *Journal of the Royal statistical society: series B (Methodological)* **57**(1), 289–300.

Benzer, T. (2022), Removing cultural barriers to education: State-run islamic schools and girls' education in turkey, Working paper.

Berndt, E., Hall, B., Hall, R. and Hausman, J. (1974), 'Estimation and inference in nonlinear structural models', *Annals of Economic and Social Measurement* **3**(4), 653–665.

Bessin, M. (2002), *Autopsie du service militaire 1965–2001*, Autrement.

Biavaschi, C., Giulietti, C. and Siddique, Z. (2017), 'The economic payoff of name americanization', *Journal of Labor Economics* **35**(4), 1089–1116.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2011), 'Ethnic identity and labour market outcomes of immigrants in europe', *Economic Policy* **26**(65), 57–92.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2016), 'Bend it like beckham: Ethnic identity and integration', *European Economic Review* **90**, 146–164.

Bisin, A. and Topa, G. (2003), 'Empirical models of cultural transmission', *Journal of the European Economic Association* **1**(2/3), 363–375.

Bisin, A., Topa, G. and Verdier, T. (2004), 'Religious intermarriage and socialization in the United States', *Journal of Political Economy* **112**(3), 615–664.

Bisin, A. and Tura, G. (2022), Marriage, fertility, and cultural integration in Italy. Working Paper.

Bisin, A. and Verdier, T. (2000), '"Beyond the Melting Pot": Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits', *The Quarterly Journal of Economics* **115**(3), 955–988.

Bisin, A. and Verdier, T. (2001), 'The economics of cultural transmission and the dynamics of preferences', *Journal of Economic Theory* **97**, 298–319.

Bisin, A. and Verdier, T. (2011), The economics of cultural transmission and socialization, *in* 'Handbook of Social Economics', Vol. 1, North-Holland, pp. 339–416.

Bisin, A. and Verdier, T. (2022), Advances in the economic theory of cultural transmission. NBER Working Paper 30466.

Bolt, U., French, E., Maccuish, J. H. and O'Dea, C. (2021), The intergenerational elasticity of earnings: exploring the mechanisms, Cepr discussion paper no. dp15975.

Botticini, M. and Eckstein, Z. (2007), 'From farmers to merchants, conversion and diaspora: Human capital and Jewish history', *Journal of the European Economic Association* **5**(5), 885–926.

Botticini, M. and Eckstein, Z. (2012), *The Chosen Few*, Princeton University Press.

Bourdieu, P. (1979), 'Les trois états du capital culturel', *Actes de la recherche en sciences sociales* **30**(1), 3–6.

Bouzar, D. and Bouzar, L. (2009), *Allah a-t-il sa place dans l'entreprise?*, Albin Michel.

Carvalho, J.-P. (2013), 'Veiling', *The Quarterly Journal of Economics* **128**(1), 337–370.

Carvalho, J.-P., Koyama, M. and Sacks, M. (2017), 'Education, identity, and community: lessons from the Jewish emancipation', *Public Choice* **171**, 119–143.

Carvalho, J.-P., Koyama, M. and Williams, C. (2022), Resisting education. Oxford Discussion paper 982.

Carvalho, J.-P. and McBride, M. (2022), The formation of religious beliefs and preferences, *in* K. Zimmerman, ed., 'Handbook of Labor, Human Resources and Population Economics', Springer.

Chaudhary, L. and Rubin, J. (2011), 'Reading, writing, and religion: Institutions and human capital formation', *Journal of Comparative Economics* **39**, 17–33.

Cheung, M.-W. and Wu, J. (2018), 'On the probabilistic transmission of continuous cultural traits', *Journal of Economic Theory* **174**, 300–323.

Chiappori, P.-A. (1992), 'Collective labor supply and welfare', *Journal of Political Economy* **100**(3), 437–467.

Chiappori, P.-A. (2017), *Matching with transfers: The economics of love and marriage*, Princeton University Press.

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

Choo, E. and Siow, A. (2006), 'Who marries whom and why', *Journal of Political Economy* **11'**(1), 175–201.

Cintas, C., Gosse, B. and Vatteville, É. (2012), Quand l'identité religieuse devient une préoccupation du management des ressources humaines, *in* 'Management et religion: Décryptage d'un lien indéfectible', EMS Editions, pp. 83–98.

Corekcioglu, G. (2021), 'Unveiling the effects of a headscarf ban: Evidence from municipal jobs in turkey', *Journal of Comparative Economics* **49**(2), 382–404.

Cunha, F. and Heckman, J. (2007), 'The technology of skill formation', *The American Economic Review* **97**(2), 31–47.

Del Boca, D., Flinn, C. and Wiswall, M. (2014), 'Household choices and child development', *The Review of Economic Studies* **81**(1), 137–185.

Del Boca, D., Flinn, C. and Wiswall, M. (2016), 'Transfers to households with children and child development', *The Economic Journal* **126**(596), F136–F183.

Del Bono, E., Francesconi, M., Kelly, Y. and Sacker, A. (2016), 'Early maternal time investments and early child outcomes', *The Economic Journal* **126**, F96–F135.

Droogsma, R. A. (2007), 'Redefining hijab: American muslim women's standpoints on veiling', *Journal of Applied Communication Research* **35**(3), 294–319.

Drydakis, N. (2013), 'The effect of ethnic identity on the employment of immigrants', *Review of Economics of the Household* **11**(2), 285–308.

Fan, X. and Wu, L. (2022), The shaping of a gender norm: Marriage, labor, and foot-binding in historical china, Working paper.

Fernández-Reino, M., Di Stasio, V. and Veit, S. (2022), 'Discrimination unveiled: A field experiment on the barriers faced by muslim women in germany, the netherlands, and spain', *European Sociological Review* .

Fouka, V. (2020), 'Backlash: The unintended effects of language prohibition in U.S. schools after World War I', *The Review of Economic Studies* **87**(1), 204–239.

Fryer, R. G. and Levitt, S. D. (2004), 'The causes and consequences of distinctively black names', *The Quarterly Journal of Economics* **119**(3), 767–805.

Galichon, A., Kominers, S. D. and Weber, S. (2019), 'Costly concessions: An empirical framework for matching with imperfectly transferable utility', *Journal of Political Economy* **127**(6), 2875–2925.

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

Garre, F. G., Vermunt, J. K. and Croon, M. A. (2002), 'Likelihood-ratio tests for order-restricted log-linear models: A comparison of asymptotic and bootstrap methods', *Metodologa de las Ciencias del Comportamiento* **4**(1), 1–18.

Gaspard, F. and Khosrokhavar, F. (1995), *Le foulard et la République*, FeniXX.

Gibbons, R. and Waldman, M. (2004), 'Task-specific human capital', *The American Economic Review* **94**(2), 203–207.

Greene, W. H. (2008), *Econometric analysis*, 6th edn, Pearson Education, Inc.

Gresh, A. (2020), 'Islam: A one-sided debate'.
URL: *https://orientxxi.info/dossiers-et-series/islam-a-one-sided-debate,4281*

Guiso, L., Sapienza, P. and Zingales, L. (2006), 'Does culture affect economic outcomes?', *Journal of Economic Perspectives* **20**(2), 23–48.

Gulesci, S., Jindani, S., La Ferrara, E., Smerdon, D., Sulaiman, M. and Young, H. (2021), A stepping stone approach to understanding harmful norms, CEPR Discussion Paper No. DP15776.

Harrison, K. A. (2016), 'Hiding under the veil of" dress policy": Muslim women, hijab, and employment discrimination in the united states', *Geo. J. Gender & L.* **17**, 831.

Heckman, J., Pinto, R. and Savelyev, P. (2013), 'Understanding the mechanisms through which an influential early childhood program boosted adult outcomes', *American Economic Review* **103**(6), 2052–86.

Hu, Z. and Wang, Z. (2021), Nutrition, labor supply, and productivity: Evidence from ramadan in indonesia, Working paper.

Iannaccone, L. R. (1990), 'Religious practice: A human capital approach', *Journal for the Scientific Study of Religion* **29**(3), 297–314.

Iannaccone, L. R. (1992), 'Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives', *Journal of Political Economy* **100**(2), 271–291.

Iannaccone, L. R. (1998), 'Introduction to the economics of religion', *Journal of Economic Literature* **36**(3), 1465–1495.

Institut Français d'Opinion Publique [IFOP] (2019), Etude auprès de la population musulmane en france, 30 ans après l'affaire des foulards de creil, Technical report.

Institut Montaigne (2014–2021), Religion au travail: croire au dialogue. baromètre du fait religieux en entreprise 2014–2021, Technical report.

Institut Montaigne (2016), A french islam is possible, Technical report.

Iyer, S. (2016), 'The new economics of religion', *Journal of Economic Literature* **54**(2), 395–441.

Jia, R. and Persson, T. (2021), 'Choosing ethnicity: The interplay between individual and social motives', *Journal of the European Economic Association* **19**(2), 1203–1248.

Jouili, J. S. (2020), *Pious Practice and Secular Constraints*, Stanford University Press.

Katz, D., Baptista, J., Azen, S. and Pike, M. (1978), 'Obtaining confidence intervals for the risk ratio in cohort studies', *Biometrics* **34**(3), 469–474.

Khosrokhavar, F. (2004), L'islam des jeunes filles en france, *in* 'Le foulard islamique en questions', Éditions Amsterdam, pp. 89–94.

Kirkebøen, L., Leuven, E. and Mogstad, M. (2022), 'College as a marriage market', *NBER Working Paper* (28688).

La Boussole (2004), *Annuaire des mosquées de France*.

Lecomte, J.-P. (2001), Représentations et réalités des fonctions sociales du service militaire dans la société française (1868–2001). Thèse de doctorat en Science politique soutenue en 2001 à Paris, Institut d'études politiques.

Livny, A. (2020), *Trust and the Islamic advantage: religious-based movements in Turkey and the Muslim world*, Cambridge University Press.

Maillard, D. (2017), *Quand la religion s'invite dans l'entreprise*, Fayard.

Maurin, E. and Navarrete-Hernandez, N. (2023), 'Behind the veil: the effect of banning the islamic veil in schools', *Economic Policy* **38**(113), 63–98.

McFadden, D. (1973), Conditional logit analysis of qualitative choice behavior, *in* P. Zarembka, ed., 'Frontiers in Econometrics', Academic Press, New York, chapter 4, pp. 105–142.

Montgomery, J. D. (2010), 'Intergenerational cultural transmission as an evolutionary game', *American Economic Journal: Microeconomics* **2**(4), 115–136.

Morsy, M. (1989), *Les femmes du Prophète*, FeniXX.

Munshi, K. (2011), 'Strength in numbers: Networks as a solution to occupational traps', *The Review of Economic Studies* **78**(3), 1069–1101.

Munshi, K. (2019), 'Caste and the Indian economy', *Journal of Economic Literature* **57**(4), 781–834.

Nordmann, C. (2004), *Le foulard islamique en questions*, Éditions Amsterdam.

Novak, L. (2020), 'Persistent norms and tipping points: The case of female genital cutting', *Journal of Economic Behavior & Organization* **177**, 433–474.

Patacchini, E. and Zenou, Y. (2016), 'Social networks and parental behavior in the intergenerational transmission of religion', *Quantitative Economics* **7**, 969–995.

Patel, D. S. (2012), 'Concealing to reveal: The informational role of islamic dress', *Rationality and Society* **24**(3), 295–323.

Platteau, J.-P. (2017), *Islam instrumentalized*, Cambridge University Press.

Read, J. G. and Bartkowski, J. P. (2000), 'To veil or not to veil? a case study of identity negotiation among muslim women in austin, texas', *Gender & society* **14**(3), 395–417.

Roth, A. E. and Sotomayor, M. A. O. (1990), *Two-sided matching: A study in game-theoretic modeling and analysis*, Cambridge University Press.

Sakalli, S. O. (2019), Secularization and religious backlash: Evidence from turkey, Technical report, Technical Report, Working Paper.

Saleh, M. (2016), 'Public mass modern education, religion, and human capital in twentieth-century Egypt', *The Journal of Economic History* **76**(3), 697–735.

Saleh, M. (2018), 'On the road to heaven: Taxation, conversions, and the Coptic–Muslim socioeconomic gap in Medieval Egypt', *The Journal of Economic History* **78**(2), 394–434.

Scott, J. W. (2009), *The politics of the veil*, Vol. 7, Princeton University Press.

Sénat (2021), 'Compte rendu intégral: Séance du mardi 30 mars 2021', *Journal Officiel de la République française* .

Shofia, N. M. (2020), Why veil? religious headscarves and the public role of women, Working paper.

Silhouette-Dercourt, V., Sy, O. S. and Desjeux, D. (2019), 'Cosmopolitan veiling in paris: Young french muslim women in transition', *Youth and Globalization* **1**(1), 65–87.

Siow, A. (2015), 'Testing Becker's theory of positive assortative matching', *Journal of Labor Economics* **33**(2), 409–441.

Skrondal, A. and Rabe-Hesketh, S. (2009), 'Prediction in multilevel generalized linear models', *Journal of the Royal Statistical Society Series A: Statistics in Society* **172**(3), 659–687.

Squicciarini, M. (2020), 'Devotion and development: Religiosity, education, and economic progress in 19th-century France', *The American Economic Review* **110**(11), 3454–3491.

Valfort, M.-A. (2020), 'Anti-muslim discrimination in france: Evidence from a field experiment', *World Development* **135**, 105022.

van der Hasselt, G. (2019), 'The muslim veil in france: Why so controversial? three questions to hakim el kaouri'.
URL: *https://www.institutmontaigne.org/en/blog/muslim-veil-france-why-so-controversial*

Weber, M. (2013), *The Protestant Ethic and the Spirit of Capitalism*, Routledge.

Weichselbaumer, D. (2020), 'Multiple discrimination against female immigrants wearing headscarves', *ILR Review* **73**(3), 600–627.

Chapter 3

Are marriage markets segmented?

3.1 Introduction

Marriage is a seminal event in a person's life, bearing significant economic implications, including on consumption, the transmission of human capital, and the persistence of inequality across generations (Becker 1973, 1974, Chiappori 1992). While we know who is getting married and when, understanding the 'why' behind marriage choices remains a pressing question in economic research (Choo and Siow 2006). Consider, for example, the rise of educational homogamy. Recent research has shown that couples in the United States are increasingly sorted by education level (Chiappori et al. 2017). This trend has been largely attributed to the growing value of human capital, as parents with similar education levels can better equip their children for future success. But could there be other, overlooked factors at play?

One such factor could be the segmentation of marriage markets, i.e. the fact that unalike individuals may be unlikely to interact, and thus to become marital partners. For instance, anecdotal evidence suggests that social diversity of students within educational institutions is shrinking, with elites increasingly keeping their children within the same institutions and social circles. This segregation could create a feedback loop of assortativeness, contributing to the rise of educational homogamy. In fact, recent work by Kirkebøen et al. (2022) shows this mechanism at play in colleges in Denmark, with college graduates finding their future spouse within the same college institution to a large extent.

In economics, the standard models for marriage are the frictionless and search models. The frictionless framework often overlooks partial market segmentation, being only able to handle cases of complete segmentation. For instance, Chiappori et al. (2017) consider distinct markets for different racial groups, but in so doing must ignore mixed-race couples. Search models, however, typically assume that search costs do not vary based on the type of individual sought.

In this paper, I study the role of market segmentation on marriage markets. Theoretically, I propose a middle-ground approach, adopting a frictionless setting within a segmented market context. Here, segmented market refers to a scenario where individuals of a given type A are more likely to encounter potential partners of type A than of type B. The model implements this idea by building on the model initially proposed by Choo and Siow (2006), but considering that individuals of different types are assigned to different submarkets where their type is overrepresented compared to the overall population. Using examples, I show how the marital assortativeness which mechanically results from this segmentation can be mistakenly interpreted as a preference for assortativeness when using standard estimates.

To provide evidence that market segmentation does play a role in marriage markets, I then turn to a quasi-natural experiment, the termination of the mandatory military service in France in 1996. Because the military service was a vector of social mixing, its termination can indeed be interpreted as a change in the structure of the marriage market. I conduct an event study around this shift to examine the impact on educational homogamy. My initial results reveal a significant shift in educational homogamy after the termination of the mandatory military service, underscoring the influence of market segmentation in marriage patterns.

Finally, in ongoing work I lay out elements for a structural estimation of the model of marriage with segmented markets. The main difficulty comes at identification: because the standard Choo–Siow model is exactly identified from marriage patterns (attributing assortativeness to preferences only), disentangling the role of market segmentation from that of preferences requires additional sources of variation. I discuss how, if identified, the model can be estimated via maximum likelihood, in which case segmentation of the market can be tested by performing a horserace between the standard Choo–Siow model and the model with segmentation.

The paper is structured as follows. Section 3.2 presents the standard Choo–Siow model, followed by the new model which introduces market segmentation. In section 3.3 I study two examples, showing the extent to which Choo–Siow estimates can misattribute the role of market segmentation in homogamy to preferences. Section 3.4 presents the event study of the termination of the mandatory military service in France. Finally, section 3.5 presents elements for a structural estimation of the model with market segmentation.

3.2 Model

The Choo and Siow (2006) model provides a practical and tractable framework to analyze matching patterns in marriage markets. It uses a frictionless matching setting with transferable utility. In this setting, men and women are categorized by distinct traits (e.g., education, race, income),

and each pair of man and woman generates a certain marital surplus which depends on their respective traits. This surplus is split between the spouses, thereby determining who marries whom. The model has been influential due to its simplicity and the fact that it can be exactly identified from observed marriage patterns in cross-sectional data.

In this section, I first present a short introduction to the Choo–Siow model with its crucial assumptions, notably the transferable utility property and the separability of errors. I then introduce market segmentation in this framework by focusing on the simple case with two categories for each spouse.

3.2.1 The Choo–Siow model

This presentation is loosely adapted from Chiappori (2017). Consider a large population of women and men on a single, non-frictional marriage market. Each individual i belongs to a category $I = 1, \dots, \mathcal{I}$. For these categories we can think for instance of educational attainment, religious affiliation, citizenship, place of residence, or any other trait. The goal of the model is to rationalize marriage patterns between women and men across these categories.

Matching with transferable utility. When a woman i and a man j marry, they produce a *surplus* which is specific to that pair of individuals, $s(i,j)$. The *transferable utility* framework assumes that this surplus can be divided between the spouses according to

$$u_i + v_j = s(i,j), \tag{3.1}$$

where u_i is the utility obtained by woman i and v_j is the utility obtained by man j.

This division of the surplus is endogenous: the utilities u_i and v_j are determined by equilibrium conditions on the marriage market. The equilibrium notion here is *stability*, which means that no two individuals, by matching together, are able to achieve greater utility than they currently have. Stability translates into the following conditions:

$$u_i = \max_j \big\{ s(i,j) - v_j \big\} \quad \text{and} \quad v_j = \max_i \big\{ s(i,j) - u_i \big\}. \tag{3.2}$$

The usual interpretation of (3.2) is that v_j is the price that any woman must "pay" to marry man j, in terms of the share of the surplus that she has to give up. She thus chooses her spouse j in order to maximize her own utility, which is the total surplus produced by the match minus the share that the man demands. Of course, a symmetric interpretation exists for u_i as the price to marry woman i.

A well-known consequence of the transferable utility property is that social surplus must be maximized at the equilibrium matching: if it weren't, we could create a new match in which both spouses would be strictly better-off, thereby violating stability. Furthermore, if the surplus function is supermodular in the spouses' traits, then perfect assortativeness arises in equilibrium. (If it is submodular, then negative assortativeness occurs.) Of course, real-world data never exhibits perfect assortativeness. For this reason, empirical models must introduce some unobserved heterogeneity to account for imperfect assortativeness.

Unobservable heterogeneity. In the Choo–Siow model, the surplus produced by a match is assumed to be the sum of a deterministic component which depends only on the categories I and J of the spouses, $S(I,J)$, and of a stochastic term ε_{ij} which is pair-specific:

$$s(i,j) = S(I,J) + \varepsilon_{ij}. \tag{3.3}$$

Furthermore, it is assumed that the stochastic term is separable into two terms:

$$\varepsilon_{ij} = \alpha_i^J + \beta_j^I. \tag{3.4}$$

These two terms α_i^J and β_j^I can be interpreted as taste heterogeneity across spouse categories for woman i and man j. Note that under this assumption, all men in category J are perfect substitutes for woman i (and vice-versa for man j). This assumption is necessary for the tractability of the model.

For singles (denoted using $\emptyset$), the deterministic part of the utility is normalized to 0:

$$s(i,\emptyset) = \alpha_i^\emptyset, \qquad s(\emptyset,j) = \beta_j^\emptyset. \tag{3.5}$$

The central result of the Choo–Siow model is that equilibrium utilities can then also be separated into a deterministic component which depends on the spouses' categories, and a stochastic component which corresponds to individuals' taste shocks. The deterministic components, U^{IJ} for women and V^{IJ} for men, verify

$$U^{IJ} + V^{IJ} = S(I,J) \qquad (\forall I, J) \tag{3.6}$$

i.e. they add up to the deterministic part of the surplus. Furthermore, if $i \in I$ is married to

$j \in J$ in equilibrium, their utilities are

$$u_i = U^{IJ} + \alpha_i^J \quad \text{and} \quad v_j = V^{IJ} + \beta_j^I. \tag{3.7}$$

As a consequence of stability, woman i is matched with a man in J in equilibrium if and only if

$$U^{IJ} + \alpha_i^J \geq U^{IL} + \alpha_i^L, \qquad \forall L \tag{3.8}$$

and she stays single if

$$\alpha_i^\emptyset \geq U^{IL} + \alpha_i^L, \qquad \forall L. \tag{3.9}$$

This is a discrete choice problem, and with a standard distributional assumption on the stochastic terms α_i^J (specifically, they are i.i.d. Gumbel) we recover a standard logistic model. The probability for a woman in I to be matched with a man in J is then

$$\begin{aligned} \gamma^{IJ} &= \Pr(\text{woman } i \in I \text{ matched with a man in } J) \\ &= \frac{\exp U^{IJ}}{1 + \sum_L \exp U^{IL}}, \end{aligned} \tag{3.10}$$

and the probability that she remains single is $\gamma^{I\emptyset} = \frac{1}{1+\sum_L \exp U^{IL}}$.

Similarly, with a corresponding hypothesis on the terms β_j^I we obtain the probability for a man in J to be matched with a woman in I as

$$\begin{aligned} \delta^{IJ} &= \Pr(\text{man } j \in J \text{ matched with a woman in } I) \\ &= \frac{\exp V^{IJ}}{1 + \sum_K \exp V^{KJ}}, \end{aligned} \tag{3.11}$$

and he remains single with probability $\delta^{\emptyset J} = \frac{1}{1+\sum_K \exp V^{KJ}}$.

This model is exactly identified from matching patterns, since the formulas (3.10) and (3.11) can be inverted to obtain the utilities U^{IJ} and V^{IJ} from the matching probabilities γ^{IJ} and δ^{IJ}. Specifically,

$$U^{IJ} = \ln \frac{\gamma^{IJ}}{\gamma^{I\emptyset}} \quad \text{and} \quad V^{IJ} = \ln \frac{\delta^{IJ}}{\delta^{\emptyset J}}. \tag{3.12}$$

Equilibrium utilities are thus estimated from empirical marriage frequencies.

3.2.2 Matching in a segmented marriage market

Now let's introduce segmentation in the Choo–Siow model. Instead of a single marriage market where any woman is allowed to match with any man in the population – and vice-versa – I will

assume that there are different marriage submarkets. Matching occurs freely *within*, but not across, these submarkets. Individuals are randomly assigned to a submarket, but in such a way that individuals of the same category are *more likely* to find themselves on the same submarket. (This is a modelling choice, because the model is built to explain positive assortativeness.) Crucially, the surplus produced by a match is independent of which submarket the match occurs on – only the division of this surplus between the spouses will be affected, reflecting the equilibrium conditions on each submarket.

This segmentation creates an alternative – possibly complementary – explanation for marriage assortativeness. In this model, assortativeness can be the result of preferences regarding spouses' traits (which takes the form of supermodularity in the surplus function), as in the Choo–Siow model, or of market segmentation on those traits.

For simplicity in the exposition, in this paper I consider what is arguably the simplest possible case to introduce segmentation. I consider only two categories of individuals: these could be college graduates and non-college graduates, Protestants and Catholics, urban and rural dwellers, etc. In general, I label these categories 1 and 2. The total number of women on the market is n, and the total number of men is m. For any category $I = 1, 2$, the number of women I is n_I, such that $n_1 + n_2 = n$. Similarly, the number of men J is m_J, with $m_1 + m_2 = m$.

Segmented market. In the Choo–Siow model, there is a single marriage market where any woman is allowed to match with any man in the population, and vice-versa. Here, I will assume instead that there are two distinct marriage markets, and that matches cannot occur across these markets. These two marriage markets are populated as follows. Individuals from category 1 always belong to market 1. On the other hand, individuals from category 2 may enter either market 1 or 2 according to exogenous probabilities. (At the end of this section I discuss how the decision to enter each market could be made endogenous.) Women 2 enter market 2 with probability ρ, and market 1 with probability $1 - \rho$. Men 2 enter market 2 with probability τ, and market 1 with probability $1 - \tau$. This exogenous market assignment is represented on Figure 3.1. Together, the probabilities ρ and τ measure the segmentation of the market. For instance, if $\rho = \tau = 0$ then everyone enters market 1 and there is no market segmentation. At the other extreme, if $\rho = \tau = 1$ then there is complete segregation of categories and the marriage market is fully segmented.

In order to interpret this market assignment mechanism, consider the example of college education. Suppose that category 1 represents non-college graduates and category 2, college graduates. College graduates may have access to a specific marriage marketplace, e.g. university

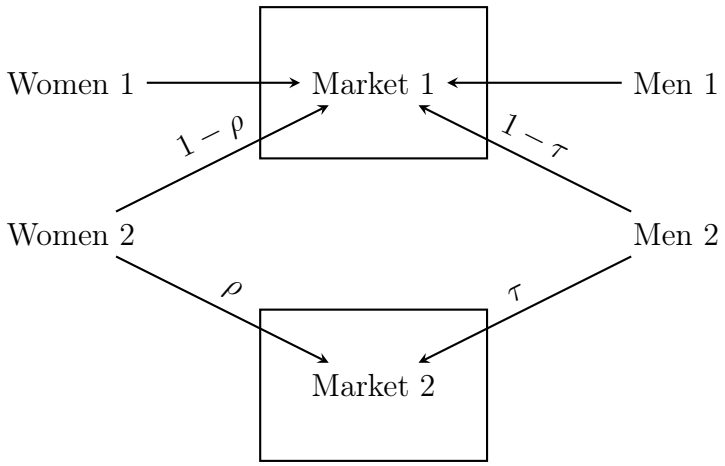

Figure 3.1: Population and marriage markets

Non-segmented market

	Women	Men
Category 1	n_1	m_1
Category 2	n_2	m_2

Segmented market

	Market 1		Market 2	
	Women	Men	Women	Men
Category 1	n_1	m_1	0	0
Category 2	$(1-\rho)n_2$	$(1-\tau)m_2$	$\rho\, n_2$	$\tau\, m_2$

Table 3.1: Composition of non-segmented versus segmented markets.

campuses or group activities for students, to which non-college graduates do not have access. Yet not all college graduates may actively look for a partner on this restricted marketplace; they may also look for a partner on the "regular" market to which non-college graduates also have access, such as coffee shops, gym classes, or community social events. In this case, the segmentation of the market represents the extent to which college graduates spend time in, and look for partners during, the activities which are restricted to them.

With the assignment mechanism described above it is easy to compute the composition of each submarket: on market 1 there are n_1 women 1, $(1-\rho)n_2$ women 2, etc. I summarize this composition in Table 3.1. This assignment structure is sufficient for both the sex ratio and the category ratio to vary with the segmentation parameters ρ and τ. For instance, in market 1, the sex ratio is equal to $\frac{n_1+(1-\rho)n_2}{m_1+(1-\tau)m_2}$, and the category ratio of women 1 to women 2 is $\frac{n_1}{(1-\rho)n_2}$.

Matching. The exogenous assignment of individuals described above creates two different marriage markets. Within these markets, I assume that matching occurs exactly as described by the Choo–Siow model. Therefore, for every spouse categories I and J, there exist parameters

U_1^{IJ} such that

$$\begin{aligned}\gamma_1^{IJ} &= \Pr(\text{woman } i \in I \text{ is matched with a man in } J \mid i \text{ in market 1}) \\ &= \frac{\exp U_1^{IJ}}{1 + \sum_L \exp U_1^{IL}},\end{aligned} \tag{3.13}$$

and for women 2 there also exists a parameter U_2 such that

$$\begin{aligned}\gamma_2 &= \Pr(\text{woman } i \in 2 \text{ is matched with a man in } 2 \mid i \text{ in market 2}) \\ &= \frac{\exp U_2}{1 + \exp U_2}.\end{aligned} \tag{3.14}$$

In these notations, the subscript refers to the market. The parameters U_1^{IJ} correspond to the equilibrium utilities obtained by matched women on market 1, while the parameter U_2 corresponds to the equilibrium utility of matched category-2 women in market 2. (We do not need to use the superscripted γ_2^{IJ} or U_2^{IJ} because there are only category-2 individuals in market 2, making the notations γ_2 and U_2 instead of γ_2^{22} and U_2^{22} unambiguous.) As before, the probability to remain single in each market is obtained by replacing the numerator by a 1 in equations (3.13) and (3.14).

Similarly for men, there are parameters V_1^{IJ} and V_2 which are linked to the matching probabilities δ_1^{IJ} and δ_2 according to corresponding formulas. Together, these parameters must satisfy the constraint

$$U_1^{22} + V_1^{22} = U_2 + V_2, \tag{3.15}$$

which states that the surplus produced by a match between category-2 spouses is the same whether they meet in market 1 or in market 2. In other words, it is assumed that market segmentation has no effect on the surplus which is produced by spouses, only on the way that this surplus is divided between the spouses. Importantly, all these equilibrium parameters ultimately depend on the probabilities ρ and τ. Indeed, the segmented structure of the marriage market will affect the utility that each individual can obtain in equilibrium.

We can then compute the ex-ante probability (before being assigned to a market) for category-2 women or men to be matched with a spouse who is also in category 2:

$$\tilde{\gamma}^{22} = (1 - \rho)\, \gamma_1^{22} + \rho\, \gamma_2 \tag{3.16}$$

$$\tilde{\delta}^{22} = (1 - \tau)\, \delta_1^{22} + \tau\, \delta_2. \tag{3.17}$$

For instance, a category-2 woman enters market 1 with probability $1 - \rho$, in which case she

matches with a category-2 man with probability γ_1^{22}; and she enters market 2 with probability ρ, in which case she matches with a category-2 man with probability γ_2. These ex-ante probabilities are important because, in the absence of information on the actual segmentation of the market, they are the matching probabilities that we will observe in the data – as opposed to the matching probabilities on each market.

The price of market segmentation. If the parameters of the model with segmentation are known, we can quantify the effect of market segmentation on spouses' welfare. Indeed, in this case we know the match surplus $S(I,J) = U_1^{IJ} + V_1^{IJ}$ for any spouse categories IJ, which by assumption is the same whether the market is segmented or not. We can therefore use standard algorithms to compute the equilibrium matching as it would occur if the market weren't segmented. By doing so we obtain the division of the surplus which would occur in equilibrium in a non-segmented market, U_0^{IJ} and V_0^{IJ}, as well as the resulting matching pattern.

With these new equilibrium utilities we can for instance quantify the average utility gain or loss of a category-2 woman due to segmentation: it is her expected utility when the market is segmented, minus her expected utility if there were no segmentation:

$$\begin{aligned}\Delta u_2 &= \underbrace{(1-\rho)\,\mathbf{E}(u_i \mid i \text{ in market } 1) + \rho\,\mathbf{E}(u_i \mid i \text{ in market } 2)}_{\text{expected utility when the market is segmented}} - \underbrace{\mathbf{E}(u_i \mid \text{no segm.})}_{\text{when not segmented}} \\ &= (1-\rho)\ln\left(1+\sum_J \exp U_1^{2J}\right) + \rho\ln\left(1+\exp U_2\right) - \ln\left(1+\sum_J \exp U_0^{2J}\right) \qquad (3.18)\end{aligned}$$

where the second equality uses a property of the Gumbel distribution to compute the expectation of the max. Similarly for category-1 women:

$$\Delta u_1 = \ln\left(1+\sum_J \exp U_1^{1J}\right) - \ln\left(1+\sum_J \exp U_0^{1J}\right). \qquad (3.19)$$

Corresponding formulas hold for men of each category as well.

Alternatively, we could also study the effect of a marginal increase or decrease in the segmentation parameters ρ and τ. To do this, we would determine the composition of the markets 1 and 2 under a small change in one of these parameters. We could then compute the equilibrium matching in each of these markets, and obtain formulas equivalent to (3.18) and (3.19) in the case of marginal changes. This computation would provide a measure of the "shadow price" of market segmentation.

Endogenizing market segmentation. An obvious limitation of the model as presented in this section is the fact that assignment to different submarkets follows exogenous probabilities. Instead, we could expect that individuals make a deliberate decision to enter the restricted market or the unrestricted one, depending on the level of utility that they expect to obtain on each market. Endogenizing the decision to enter the restricted market would also bring the model closer to the structured marriage market described in Bisin and Verdier (2000), who studied assortativeness on cultural traits.

This endogenous decision would add a new layer of complexity to the problem, in large part because sorting based on individual taste shocks would occur as a result. Indeed, imagine that a category-2 man has a strong idiosyncratic preference for women 1. That individual would have a relatively higher incentive to enter the unrestricted market, compared to a category-2 man with average preference for women 1. From this simple observation, we can infer that there would be some threshold in the value of this idiosyncratic preference, such that men 2 above that threshold enter market 1, and men 2 below it enter market 2. (Actually, this threshold would also depend on the idiosyncratic preference for women 2, but let's ignore that for the sake of the argument.)

In consequence, the distribution of taste shocks among men 2 in each submarket would be skewed. Because they are rational, women would realize that, and this knowledge would affect the share of the surplus that they can extract from men on each of the two markets. Women 1 would be able to ask for a larger share of the surplus on market 1, and women 2 on market 2. Unlike in the Choo–Siow model, equilibrium utilities asked by each side of the market would depend on the distribution of the idiosyncratic shocks.

Finally, the reasoning that I described here for men would also occur on the other side of the market, for women. In equilibrium, each side of the market would infer the cut-off of the distribution of idiosyncratic tastes, thereby calculating how much utility they can expect to claim on each market.

Overall, endogenizing the segmentation of the market is a desirable addition to this model, although not a completely straightforward one. This is a promising avenue for future work.

3.3 Examples and model misspecification

In order to better understand the impact of market segmentation on the standard Choo–Siow framework, in this section I perform a couple of exercises. In two different examples, I consider a surplus function and a segmented market as described in section 3.2, and I compute the equilibrium matching in this segmented market. I then estimate the surplus function from aggregate matching patterns, as one would do if they ignored the underlying segmented structure

of the market. In both examples, I find that the estimated surplus significantly departs from the actual postulated surplus function: as expected, the Choo–Siow framework rationalizes the segmentation-driven assortativeness by an excess complementarity on the spouses' traits. Finally, I provide estimates of the cost of market segmentation on individuals' utilities.

Example 1. In order to demonstrate how market segmentation can be interpreted as a preference for assortativeness in the Choo–Siow model, let us start with a very simple model with segmentation, but without complementarity on traits. My goal here is to measure how Choo–Siow estimates of the surplus function depart from the actual surplus function because of the market segmentation.

In this simple example, I consider the overall marriage market to be fully symmetric. There are as many women as there are men, say $n = m = 1000$ for illustration, with 40% of both women and men being college graduates. Suppose further that among women and men college graduates, $\rho = \tau = 20\%$ enter 'market 2,' the marriage submarket restricted to college graduates. That means 80% of college graduates enter 'market 1,' where they can match with non-college graduates. This is a relatively weak segmentation, with only $20\% \times 40\% \times 1000 = 80$ out of 1000 individuals of a given gender belonging to the restricted market.

Finally, suppose that there are gains to marriage compared to being single, but that those gains are independent of the spouses' types. The surplus is thus defined as

$$S = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \tag{3.20}$$

meaning that two married spouses together produce 1 more util than if they were single.

From this model, I compute the marriage patterns expected in the Choo-Siow framework (I use the method from Galichon and Salanié 2022, which consists in finding the matching equilibrium and associated values through the optimization of a regularized linear programming problem). These marriage patterns are presented in Table 3.2.

Our next step is to compute the Choo–Siow estimates of the surplus function based on the aggregate marriage patterns across the two markets, according to the formula

$$\hat{S}_{IJ} = \log \frac{(\mu_{IJ})^2}{\mu_{I\emptyset}\,\mu_{\emptyset J}} \tag{3.21}$$

where μ_{IJ} is the number of empirical matches between women I and men J, $\mu_{I\emptyset}$ the number of single women I, and $\mu_{\emptyset J}$ the number of single men J. Using the marriage patterns aggregated from both markets in Table 3.2, this formula leads to the following estimate for the surplus

		Market 1				Market 2		
			Men		Total		Men	Total
			NC	CG			NC CG	
		Singles	165	59		Singles	30	
Women	NC	165	272	163	600			
	CG	59	163	98	320	30	50	80
Total			600	320	920		80	80

Table 3.2: Example 1, expected matching patterns.

matrix:

$$\hat{S} = \begin{pmatrix} 1.00 & 0.59 \\ 0.59 & 1.00 \end{pmatrix}. \tag{3.22}$$

On the off-diagonal, this estimate is sensibly different from the actual surplus matrix (3.20), which was defined with all its elements equal to 1. As expected, market segmentation is interpreted in the Choo–Siow framework as some form of complementarity in the spouses' traits: although we did not assume any such complementarity, the estimated surplus matrix is clearly supermodular (i.e. $\hat{S}_{11} + \hat{S}_{22} > \hat{S}_{12} + \hat{S}_{21}$). Here, the off-diagonal surplus must be almost halved from its actual value in order to explain the empirical assortativeness due to the market segmentation. This is despite the relatively weak segmentation which is postulated in this example.

Finally, I can compute the cost of market segmentation, i.e. the utility gain or loss which is caused by segmentation as formulated in equations (3.18) and (3.19):

$$\Delta u_1 = \Delta v_1 = -0.05, \qquad \Delta u_2 = \Delta v_2 = -0.05. \tag{3.23}$$

All individuals on the market are losing from market segmentation. Because the market is gender-symmetric, we expect that women and men of a given category I should lose exactly as much, i.e. $\Delta u_I = \Delta v_I$. As it happens, in this example category-1 individuals lose approximately as much as category-2 individuals (this is a coincidence of this particular specification – actually, examining more decimals shows that these values are different). This -0.05 utility loss caused by segmentation represents a 4% loss for category-1 individuals, and a 3% loss for category-2 individuals, compared to the level of utility that they achieve when the market is not segmented.

Example 2. Let us replicate this exercise using a slightly more complicated example. Suppose that men outnumber women by 10%, $n = 1000$ and $m = 1100$. Additionally, 40% of men but only 20% of women are now college graduates. That means men college graduates outnumber

		Market 1				Market 2		
			Men		Total		Men	Total
			NC	CG			NC CG	
		Singles	257	22		Singles	50	
Women	NC	219	392	189	800			
	CG	0	11	9	20	10	170	180
Total			660	220			220	80

Table 3.3: Example 2, expected matching patterns.

women college graduates by a lot: 200 women versus 440 men are college graduates. For this reason, suppose that men enter the restricted market less than women do: $\rho = 90\%$ and $\tau = 50\%$, so that there are 180 women and 220 men on the restricted market.

Finally, let us assume this time that the actual surplus *is* supermodular:

$$S = \begin{pmatrix} 1 & 2 \\ 2 & 4 \end{pmatrix}, \tag{3.24}$$

so that college graduates are more valuable on the marriage market, and that there is some complementarity between spouses' traits.

As in example 1, I compute the marriage patterns predicted by the Choo–Siow framework on each of the submarkets – they are presented in Table 3.3. We note that women college graduates, who are both sought-after (because of the supermodularity of the surplus) and in short supply, attain very high marriage rates overall, compared to other categories. Men college graduates also have high marriage rates on the unrestricted market, where they are also valuable compared to their non-college graduates counterparts. Because very few college graduate women enter the unrestricted market, there are notably very few matches between women college graduates and men non-college graduates.

I now compute the Choo–Siow estimate of the surplus matrix according to the aggregated marriage patterns:

$$\hat{S} = \begin{pmatrix} 1.00 & 0.81 \\ -3.09 & 2.73 \end{pmatrix}. \tag{3.25}$$

This estimate differs significantly from the actual surplus matrix (3.24). The very low number of matches between college graduate women and non-college graduate men is rationalized through a very low surplus produced by this type of couple. Similarly, but to a lesser extent, the limited number of matches between college graduate men and non-college graduate women

is also explained by a low surplus. Of course, we know that both these patterns are to a large extent a consequence of the built-in segmentation in the model, but they are interpreted by the Choo–Siow estimates as a significant departure from actual surplus.

Again, I end this example by computing the price of segmentation for the different individuals:

$$\Delta u_1 = -0.17 \text{ (12\% loss)}, \qquad \Delta u_2 = -0.74 \text{ (19\% loss)}, \tag{3.26}$$

$$\Delta v_1 = -0.10 \text{ (9\% loss)}, \qquad \Delta v_2 = -0.19 \text{ (9\% loss)}. \tag{3.27}$$

Women 2 are the main losers from the segmentation of the market. By examining the utility levels that they obtain on the different markets, we observe that this loss mainly comes from the fact that they obtain much less utility on the restricted market (where 90% of them go), compared to if the market weren't segmented. Indeed, by entering the restricted market they reduce their bargaining power, which would be very high in the absence of segmentation because (1) there are fewer women than men in total, and (2) category-1 men also value category-2 women.

On the contrary, those who lose the least from the market segmentation (in absolute terms) are category-1 men. Indeed, they benefit from lower competition from category-2 men who enter the restricted market, increasing their bargaining power with category-1 women.

Overall, these two examples illustrate how market segmentation can lead to overestimating the importance of preferences in marriage patterns. A misspecified model which ignores the segmentation of the market produces estimates of marital preferences which depart significantly from the actual preferences. Furthermore, we saw how market segmentation can impact the equilibrium welfare of individuals, by changing the sex ratio and the category ratio on submarkets.

3.4 Market segmentation and the military service

In section 3.3 we have seen how ignoring the segmented structure of the market could lead to overestimating the role of preferences in matching patterns. In this section, we turn to the data in order to obtain evidence on whether such a structure exists in an actual marriage market.

In order to identify the role of market segmentation from that of spouses' preferences on marriage patterns, I use the termination of the mandatory military service in France. This is an interesting event to study in relation to the segmentation of the marriage market for two reasons. First, the military service was a vector of social mixing for men, since it brought together individuals from different geographical regions, socioeconomic statuses, and educational

backgrounds. This exposure to a diverse social network could influence their subsequent life choices, including partner selection. Second, the termination of the mandatory military service can be considered as a quasi-natural experiment: in 1996, it was decided that male cohorts born from 1979 onwards would not serve, so that previous cohorts were 'treated' and posterior cohorts were not. The landscape of social interactions was suddenly altered for these cohorts, and potentially the structure of the marriage market along with it.

In the following analysis I study the effect of this event on educational homogamy. To make the analysis as simple as possible, I focus on two educational levels: individuals with a tertiary education degree (e.g. a college degree), and those without.

The military service in France. The mandatory military service was compulsory in France until 1996. It lasted 18 months from 1950 until 1970, when it was reduced to 12 months (affecting the cohorts born from 1952 onwards) in a push by then-President Georges Pompidou to modernize the French military and reduce the burden of service on French youth. In 1996, President Jacques Chirac announced a professionalization plan for the French military, with the idea of transitioning to a voluntary, professional military force. That year, mandatory military service was officially suspended for all cohorts born after 1979 (those who turned 18 from 1997 onwards), as they were not called for conscription. This effectively ended the general draft in France.

The mandatory military service has been regularly described as a source of social mixing (Lecomte 2001, Bessin 2002). Indeed, 18-years-old recruits were usually affected to a military division independently of their individual characteristics (for the most part), and before they undertook their upper education, thereby limiting the possibility for educational sorting. In this setting, young men interacted beyond usual social boundaries, potentially influencing their future social networks, and as a consequence their marital decisions. The cessation of this service may have reshaped the marriage market dynamics for affected cohorts, narrowing their pool of potential partners.

Data. I use the *Trajectoires et Origines* dataset to examine marriage patterns across cohorts in France. This survey conducted in 2008 contains information on more than 20,000 individuals residing in France, born between 1948 and 1990 (i.e. between 18 and 60 years old at the time of survey). I remove first-generation immigrants from the sample, because they were not subject to the military draft. Furthermore, I combine observations into 5-year cohorts (e.g. men born between 1974 and 1978) in order not to dilute observations too much across years. This construction of 5-year cohorts leads me to consider only individuals born in the period 1949–1988,

Male cohort	1949–53	1954–58	1959–63	1964–68	
Number of obs.	303	538	857	1153	
	1969–73	1974–78	1979–83	1984–88	Total
	1374	1281	965	321	6792

Table 3.4: Number of observed matches per male cohort

representing 8 such cohorts. Finally, I remove singles (men and women) from the sample: this is because single women cannot be linked to a male cohort in a straightforward way (unlike partnered women), and therefore I leave them out of the analysis for now.[1] By symmetry, I must then also ignore single men. This brings the total number of observed matches to 6,792. Table 3.4 provides the distribution of these observations among the 5-year cohorts.

Educational homogamy across male cohorts. To study the evolution of educational homogamy across male cohorts, I use a common measure of assortativeness, the log-cross-products ratio (log-CPR) statistic. Denoting by μ_{IJ} the number of women of type I matched with a man of type J, it is defined as

$$\log \frac{\mu_{11}\,\mu_{22}}{\mu_{12}\,\mu_{21}}. \tag{3.28}$$

The log-CPR is equal to 0 when the matching pattern is consistent with random matching, positive under positive assortativeness, and negative under negative assortativeness.

In the context of the Choo–Siow model, the log-CPR is also an empirical measure of the preferences for assortativeness. To see this, consider the cross-difference obtained by summing the diagonal elements of the surplus matrix $S = (S_{IJ})_{1\leq I,J\leq 2}$ and subtracting its off-diagonal elements:

$$\mathcal{D}(S) = S_{11} + S_{22} - S_{12} - S_{21}. \tag{3.29}$$

The number $\mathcal{D}(S)$ is called the *supermodular core* of the surplus matrix S (Chiappori et al. 2017). An increase in the supermodular core means that the surplus produced by partners of the same type increases relative to the surplus produced by partners of different types. As such, it denotes an increase in the preferences for assortativeness. Using formula (3.21) for the estimate of S_{IJ},

[1] A possible way to include single women in the analysis would be to compute the empirical probability, conditional on having a partner, for a woman born in year y that her partner belongs to cohort c. Single women can then be assigned based on their year of birth to male cohorts according to those probabilities.

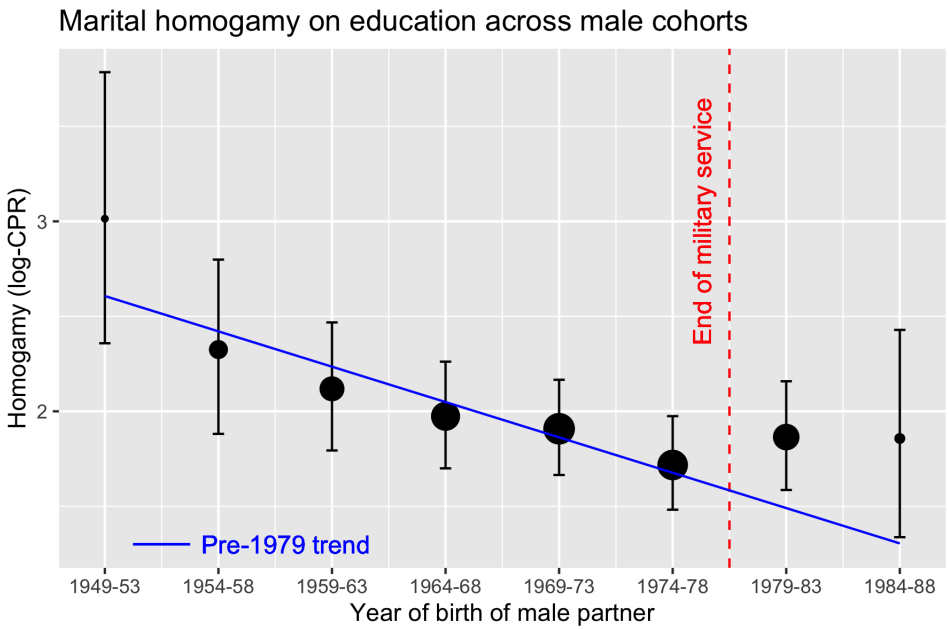

Figure 3.2: Educational homogamy and the end of the military service. Homogamy is measured using the log-CPR statistic. The size of each data point is proportional to the number of observations for that cohort (see Table 3.4). 95% confidence intervals for each point are constructed by parametric bootstrap. The pre-1979 trend (in blue) is obtained by weighted OLS, with weights equal to the number of observations at each point.

we see that the log-CPR is in fact an estimate (up to a factor 2) of the supermodular core $\mathcal{D}(S)$:

$$\begin{aligned} 2\log\frac{\mu_{11}\,\mu_{22}}{\mu_{12}\,\mu_{21}} &= \log\left(\frac{(\mu_{11}\,\mu_{22})^2}{\mu_{1\emptyset}\,\mu_{\emptyset 1}\;\mu_{2\emptyset}\,\mu_{\emptyset 2}}\;\frac{\mu_{1\emptyset}\,\mu_{\emptyset 2}\;\mu_{2\emptyset}\,\mu_{\emptyset 1}}{(\mu_{12}\,\mu_{21})^2}\right) \\ &= \log\left(\frac{(\mu_{11})^2}{\mu_{1\emptyset}\,\mu_{\emptyset 1}}\right) + \log\left(\frac{(\mu_{22})^2}{\mu_{2\emptyset}\,\mu_{\emptyset 2}}\right) - \log\left(\frac{(\mu_{12})^2}{\mu_{1\emptyset}\,\mu_{\emptyset 2}}\right) - \log\left(\frac{(\mu_{21})^2}{\mu_{2\emptyset}\,\mu_{\emptyset 1}}\right) \\ &= \hat{S}_{11} + \hat{S}_{22} - \hat{S}_{12} - \hat{S}_{21} = \widehat{\mathcal{D}(S)}. \end{aligned}$$

Thus, an increase in the log-CPR in the data should be interpreted as an increase in the preferences for assortativeness (and similarly for a decrease).

I compute the log-CPR for each of the male cohorts introduced above, along with 95% confidence intervals obtained by parametric bootstrap, and plot the results in Figure 3.2. We observe a strong positive assortativeness on education overall, as measured by log-CPR estimates which are significantly above 0. Yet, we also notice that educational homogamy has actually decreased on the period considered, as evidenced by a downwards trend in the log-CPR statistic.

This downwards trend however seems to subside, or even reverse, for the cohorts who were not subject to the mandatory military service. This is confirmed by computing the pre-1979 trend of the log-CPR statistic (blue in Figure 3.2). The post-1979 cohorts significantly depart from this trend, providing evidence of a change in trend around the cohorts affected by the policy

change.

This simple analysis provides some preliminary evidence for the role of market structure in marriage patterns. Indeed, a possible interpretation of the reversal in the trend of homogamy is that with the end of the military service, the marriage market suddenly got more segmented. The increase in assortativeness on education which followed the termination of the service was not due to a change in preferences, but rather to a change in the structure of the social interactions which can lead to finding a partner. Although it is possible that a gradual change of preferences also occurred around the same time, the fact that the discontinuity in the evolution of educational homogamy coincides with the termination of the mandatory service may suggest a causal link.

More data is needed to study the role of the termination of the mandatory military service on the structure of the marriage market. The second wave of the *Trajectoires et Origines* survey, conducted in 2018 and made available recently, could prove a useful complement to take this analysis further. Another potential source of data is the *Echantillon Demographique Permanent*, which is a 1/50 extract of the French census (containing all individuals born on the first week of October). As it stands, this simple analysis provides some basis for the argument that the segmentation of the market is an important factor in educational homogamy. Hence it is consistent with the results of Kirkebøen et al. (2022) on educational homogamy in Denmark.

3.5 Structural approach

In this section I lay out a different approach to measure the role of market segmentation in marriage patterns. The goal is to estimate the structural model of matching on a segmented market developed in section 3.2. Indeed, estimating the utility parameters along with the market segmentation parameters would allow for direct comparison of the role of preferences versus market segmentation in marriage patterns. I show that the model can be estimated via maximum likelihood, provided its parameters are identified, and I discuss possible avenues for identification.

3.5.1 General framework

To investigate whether marriage market segmentation occurs in the data, one possibility is to run a horserace between the Choo–Siow model and the model with a segmented market developed in section 3.2. The equilibrium utilities U^{IJ} and V^{IJ} of the Choo–Siow model are exactly identified from marriage patterns, and for that reason they are usually estimated directly from simple computation of (3.12).

Still, we can form the log-likelihood associated with the Choo–Siow model:

$$\ln L = \sum_{\text{women } i} \left[\sum_{I,J} \mathbf{1}_{\{i \in I \text{ matched with a man } J\}} \ln \gamma^{IJ} + \sum_{I} \mathbf{1}_{\{i \in I \text{ single}\}} \ln \gamma^{I\emptyset} \right] + \sum_{\text{men } j} \left[\sum_{I,J} \mathbf{1}_{\{j \in J \text{ matched with a woman } I\}} \ln \delta^{IJ} + \sum_{I} \mathbf{1}_{\{j \in J \text{ single}\}} \ln \delta^{\emptyset J} \right] \tag{3.30}$$

where the probabilities γ^{IJ} and δ^{IJ} depend on the parameters U^{IJ} and V^{IJ} according to equations (3.10) and (3.11). Denoting by μ_{IJ} the number of empirical matches between women I and men J, $\mu_{I\emptyset}$ the number of single women I, and $\mu_{\emptyset J}$ the number of single men J, this expression simplifies to

$$\ln L = \sum_{I,J} \mu_{IJ} \left(\ln \gamma^{IJ} + \ln \delta^{IJ} \right) + \sum_{I} \mu_{I\emptyset} \ln \gamma^{I\emptyset} + \sum_{J} \mu_{\emptyset J} \ln \delta^{\emptyset J}. \tag{3.31}$$

Differentiation of this expression yields the maximum likelihood estimates of the parameters U^{IJ} and V^{IJ}. The first-order conditions associated with these parameters are respectively

$$\frac{\mu_{IJ}}{n_I} = \frac{\exp U^{IJ}}{1 + \sum_L \exp U^{IL}} \tag{3.32}$$

$$\frac{\mu_{IJ}}{m_J} = \frac{\exp V^{IJ}}{1 + \sum_K \exp V^{KJ}}. \tag{3.33}$$

They are the empirical equivalents of the model equations (3.10) and (3.11), since μ_{IJ}/n_I is the proportion of women I matched with a man J, and μ_{IJ}/m_J is the proportion of men J matched with a woman I. This confirms that the maximum likelihood estimators of U^{IJ} and V^{IJ} coincide with the estimates obtained by direct computation from inverting (3.10) and (3.11). Thus the maximum likelihood approach is equivalent to the usual estimation by direct computation.

Now let us consider the model with segmentation. Compared to the Choo–Siow model, the number of parameters to estimate increases: there are now the equilibrium utilities of women and men on market 1, U_1^{IJ} and V_1^{IJ}, the utilities of category-2 women and men on market 2, U_2 and V_2, and the segmentation parameters ρ and τ. In total, this makes 4 additional parameters compared to the Choo–Siow model, although the constraint (3.15) (which stipulates that suplus is the same whether individuals meet on market 1 or market 2) reduces the number of additional free parameters to 3.

Unlike in the Choo–Siow model however, the formulas (3.16) and (3.17) of the model with segmentation cannot be inverted to directly obtain estimates of the model parameters. Indeed,

there are now more parameters than moments based on matching patterns – we will discuss this identification issue later. Despite this, it is still possible to write the log-likelihood of the model:

$$\ln \tilde{L} = \sum_{I,J} \mu_{IJ} \left(\ln \tilde{\gamma}^{IJ} + \ln \tilde{\delta}^{IJ} \right) + \sum_{I} \mu_{I\emptyset} \ln \tilde{\gamma}^{I\emptyset} + \sum_{J} \mu_{\emptyset J} \ln \tilde{\delta}^{\emptyset J} \tag{3.34}$$

where the probabilities $\tilde{\gamma}^{IJ}$ and $\tilde{\delta}^{IJ}$ are the ex-ante matching probabilities, which depend on the model parameters according to equations (3.16) and (3.17). Therefore, the parameters of the model can still be estimated by maximum likelihood, provided that we have enough identifying variation.

Finally, how do we test the hypothesis of market segmentation? As I explained in section 3.2, the Choo–Siow model is nested in the segmented market model. In fact, it corresponds to the following parameter values:

$$\mathbf{H_0}: \ \rho = \tau = 0. \tag{3.35}$$

Since non-segmentation has been the standard assumption in the literature, it is natural to consider $\mathbf{H_0}$ as the null hypothesis.

Because the null hypothesis corresponds to a model which is nested in the broader class of models with a segmented market, the tool of choice to detect segmentation (i.e. to test whether $\mathbf{H_0}$ holds or not) is the likelihood-ratio test. Assuming that we are able to identify the parameters of the model with segmentation, then we would be able to perform a horserace between the Choo–Siow model and the model with segmentation by comparing their maximum log-likelihoods, i.e. their ability to explain the data.

3.5.2 Identification

As it has been mentioned several times already, the Choo–Siow model is exactly identified from marriage patterns. This means that the model with a segmented market, with 3 additional free parameters, is under-identified from marriage patterns alone.

Although I do not have a definitive solution to this identification problem yet, in this last section I lay out some preliminary thoughts on potential avenues to solve it. The basic observation is that in order to identify the role of market segmentation versus that of changes in preferences, one has to rely on exogenous variation in the former or in the latter. The analysis of the end of the mandatory military service performed in section 3.4 provides some indication that exogenous variation in the structure of the matching market could be used as a source of identification.

Parametric assumptions. A possible option to reach identification is to parametrize the model in such a way as to attribute sudden changes in the estimates of the surplus to a change in the segmentation of the market. This could be a reasonable approach if we expect, as in the example of section 3.4, that the segmentation of the market changed might have changed abruptly from one cross-section to the next.

A possible approach is to impose a trend structure on the marital surplus, instead of allowing for free estimation of this surplus and individual utilities for every single cross-section. Such an assumption provides both identifying power (because departures from this trend would be interpreted as changes in the structure of the market) as well as statistical power (because the number of parameters representing preferences would decrease as long as the time horizon is long enough). With a quadratic trend for instance, the surplus of a couple IJ for year t would take the form

$$S_t^{IJ} = a^{IJ} + b^{IJ}\, t + c^{IJ}\, t^2. \tag{3.36}$$

As long as the number of periods is greater than 4, this assumption reduces the number of free parameters of the model. By constraining the shape of the surplus, the model loses in flexibility. However, this loss comes with an additional structure which allows us to analyze whether the residual variation in surplus is correlated with modifications to the segmentation of the market. Furthermore, a quadratic trend seems a reasonable approximation given what we know about the evolution of marital surplus on education in the US (Chiappori et al. 2017), or even with the empirical analysis presented in section 3.4.

3.6 Conclusion

This paper presented an alternative perspective on marriage markets, examining market segmentation and its implications for our understanding of marital homogamy. By building upon the Choo–Siow model, I considered a scenario in which individuals of a given type are more likely to encounter potential partners of the same type, as they are assigned to submarkets where their type is overrepresented. I examined how this segmentation would bias the usual Choo–Siow estimates of preferences for assortativeness, as well as the cost of this segmentation for individuals on the marriage market.

The theoretical approach was supplemented with empirical evidence of market segmentation, using a quasi-natural experiment: the termination of mandatory military service in France in 1996. This event served as a significant disruption to social mixing, thus providing a unique opportunity to observe its effects on the structure of the marriage market. The results of a

preliminary event study revealed a significant shift in educational homogamy following the end of the military service, supporting the hypothesis that market segmentation does influence marriage patterns.

Finally, I proposed elements for a structural estimation of the model of marriage with market segmentation, showing that structural estimation can be achieved via maximum likelihood, although this requires a yet-unspecified identification strategy. If identified, this approach could pave the way for further testing and refinement of the model, offering more nuanced insights into the mechanisms driving marital homogamy.

While the empirical exercise in this paper has focused on educational homogamy, other spousal categories could be studied to disentangle the role of preferences from that of market segmentation in marriage patterns. Further work could examine for instance the segmentation of the marriage market based on race, immigration status, religion, etc. In particular, historical events which may have led to a sudden change in the segmentation of the market along one of these dimensions should be studied, as they represent both an opportunity to study the role of market segmentation, as well as a source of identification for the structural estimation of the model presented here.

Further work should also seek to account for the endogeneity of individuals' decision to participate in a restricted marriage market, versus a mixed one. This consideration would establish a bridge between search models and frictionless models of marriage, in addition to providing an economic rationale behind the segmentation of the market.

Bibliography

Abdelgadir, A. and Fouka, V. (2020), 'Political secularism and Muslim integration in the West: Assessing the effects of the French Headscarf Ban', *American Political Science Review* **114**(3), 707–723.

Abdelhadi, E. (2019), 'The hijab and muslim women's employment in the united states', *Research in Social Stratification and Mobility* **61**, 26–37.

Abramitzky, R., Boustan, L. and Eriksson, K. (2020), 'Do immigrants assimilate more slowly today than in the past?', *American Economic Review: Insights* **2**(1), 125–41.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2010), 'Identifying barriers to muslim integration in france', *Proceedings of the National Academy of Sciences* **107**(52), 22384–22390.

Adida, C. L., Laitin, D. D. and Valfort, M.-A. (2016), *Why Muslim integration fails in Christian-heritage societies*, Harvard University Press.

Aizer, A. and Cunha, F. (2012), The production of human capital: Endowments, investments and fertility. NBER Working Paper 18429.

Ajbli, F. (2011), Les Françaises musulmanes face à l'emploi: le cas des pratiquantes "voilées" dans la métropole lilloise, PhD thesis, Paris, EHESS.

Akerlof, G. A. and Kranton, R. E. (2000), 'Economics and identity', *The quarterly journal of economics* **115**(3), 715–753.

Aksoy, O. (2017), 'Motherhood, sex of the offspring, and religious signaling', *Sociological Science* **4**, 511–527.

Aksoy, O. and Gambetta, D. (2016), 'Behind the veil: the strategic use of religious garb', *European Sociological Review* **32**(6), 792–806.

Aksoy, O. and Gambetta, D. (2021), 'The politics behind the veil', *European Sociological Review* **37**(1), 67–88.

Algan, Y., Malgouyres, C., Mayer, T. and Thoenig, M. (2022), 'The economic incentives of cultural transmission: Spatial evidence from naming patterns across france', *The Economic Journal* **132**(642), 437–470.

Andriantsimbazovina, J., Kabou, P. et al. (2020), *Laïcité et défense de l'État de droit*, Presses de l'Université Toulouse 1 Capitole.

Atasoy, Y. (2006), 'Governing women's morality: a study of islamic veiling in Canada', *European Journal of Cultural Studies* **9**(2), 203–221.

Atkin, D., Colson-Sihra, E. and Shayo, M. (2021), 'How do we choose our identity? a revealed preference approach using food consumption', *Journal of Political Economy* **129**(4), 1193–1251.

Beauchemin, C., Hamel, C. and Simon, P. (2016), *Trajectoires et origines: enquête sur la diversité des populations en France*, INED éditions.

Becker, G. S. (1965), 'A theory of the allocation of time', *The Economic Journal* **75**(299), 493–517.

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

Becker, G. S. (1974), 'A theory of marriage: Part II', *Journal of Political Economy* **82**(2), S11–S26.

Becker, S. O., Grosfeld, I., Grosjean, P., Voigtländer, N. and Zhuravskaya, E. (2020), 'Forced migration and human capital: Evidence from post-WWII population transfers', *The American Economic Review* **110**(5), 1430–63.

Becker, S. O. and Woessmann, L. (2009), 'Was Weber wrong? A human capital theory of Protestant economic history', *The Quarterly Journal of Economics* **124**(2), 531–596.

Bellemare, M. F., Novak, L. and Steinmetz, T. L. (2015), 'All in the family: Explaining the persistence of female genital cutting in west africa', *Journal of Development Economics* **116**, 252–265.

Benjamini, Y. and Hochberg, Y. (1995), 'Controlling the false discovery rate: a practical and powerful approach to multiple testing', *Journal of the Royal statistical society: series B (Methodological)* **57**(1), 289–300.

Benzer, T. (2022), Removing cultural barriers to education: State-run islamic schools and girls' education in turkey, Working paper.

Berndt, E., Hall, B., Hall, R. and Hausman, J. (1974), 'Estimation and inference in nonlinear structural models', *Annals of Economic and Social Measurement* **3**(4), 653–665.

Bessin, M. (2002), *Autopsie du service militaire 1965–2001*, Autrement.

Biavaschi, C., Giulietti, C. and Siddique, Z. (2017), 'The economic payoff of name americanization', *Journal of Labor Economics* **35**(4), 1089–1116.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2011), 'Ethnic identity and labour market outcomes of immigrants in europe', *Economic Policy* **26**(65), 57–92.

Bisin, A., Patacchini, E., Verdier, T. and Zenou, Y. (2016), 'Bend it like beckham: Ethnic identity and integration', *European Economic Review* **90**, 146–164.

Bisin, A. and Topa, G. (2003), 'Empirical models of cultural transmission', *Journal of the European Economic Association* **1**(2/3), 363–375.

Bisin, A., Topa, G. and Verdier, T. (2004), 'Religious intermarriage and socialization in the United States', *Journal of Political Economy* **112**(3), 615–664.

Bisin, A. and Tura, G. (2022), Marriage, fertility, and cultural integration in Italy. Working Paper.

Bisin, A. and Verdier, T. (2000), '"Beyond the Melting Pot": Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits', *The Quarterly Journal of Economics* **115**(3), 955–988.

Bisin, A. and Verdier, T. (2001), 'The economics of cultural transmission and the dynamics of preferences', *Journal of Economic Theory* **97**, 298–319.

Bisin, A. and Verdier, T. (2011), The economics of cultural transmission and socialization, *in* 'Handbook of Social Economics', Vol. 1, North-Holland, pp. 339–416.

Bisin, A. and Verdier, T. (2022), Advances in the economic theory of cultural transmission. NBER Working Paper 30466.

Bolt, U., French, E., Maccuish, J. H. and O'Dea, C. (2021), The intergenerational elasticity of earnings: exploring the mechanisms, Cepr discussion paper no. dp15975.

Botticini, M. and Eckstein, Z. (2007), 'From farmers to merchants, conversion and diaspora: Human capital and Jewish history', *Journal of the European Economic Association* **5**(5), 885–926.

Botticini, M. and Eckstein, Z. (2012), *The Chosen Few*, Princeton University Press.

Bourdieu, P. (1979), 'Les trois états du capital culturel', *Actes de la recherche en sciences sociales* **30**(1), 3–6.

Bouzar, D. and Bouzar, L. (2009), *Allah a-t-il sa place dans l'entreprise?*, Albin Michel.

Carvalho, J.-P. (2013), 'Veiling', *The Quarterly Journal of Economics* **128**(1), 337–370.

Carvalho, J.-P., Koyama, M. and Sacks, M. (2017), 'Education, identity, and community: lessons from the Jewish emancipation', *Public Choice* **171**, 119–143.

Carvalho, J.-P., Koyama, M. and Williams, C. (2022), Resisting education. Oxford Discussion paper 982.

Carvalho, J.-P. and McBride, M. (2022), The formation of religious beliefs and preferences, *in* K. Zimmerman, ed., 'Handbook of Labor, Human Resources and Population Economics', Springer.

Chaudhary, L. and Rubin, J. (2011), 'Reading, writing, and religion: Institutions and human capital formation', *Journal of Comparative Economics* **39**, 17–33.

Cheung, M.-W. and Wu, J. (2018), 'On the probabilistic transmission of continuous cultural traits', *Journal of Economic Theory* **174**, 300–323.

Chiappori, P.-A. (1992), 'Collective labor supply and welfare', *Journal of Political Economy* **100**(3), 437–467.

Chiappori, P.-A. (2017), *Matching with transfers: The economics of love and marriage*, Princeton University Press.

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

Choo, E. and Siow, A. (2006), 'Who marries whom and why', *Journal of Political Economy* **11'**(1), 175–201.

Cintas, C., Gosse, B. and Vatteville, É. (2012), Quand l'identité religieuse devient une préoccupation du management des ressources humaines, *in* 'Management et religion: Décryptage d'un lien indéfectible', EMS Editions, pp. 83–98.

Corekcioglu, G. (2021), 'Unveiling the effects of a headscarf ban: Evidence from municipal jobs in turkey', *Journal of Comparative Economics* **49**(2), 382–404.

Cunha, F. and Heckman, J. (2007), 'The technology of skill formation', *The American Economic Review* **97**(2), 31–47.

Del Boca, D., Flinn, C. and Wiswall, M. (2014), 'Household choices and child development', *The Review of Economic Studies* **81**(1), 137–185.

Del Boca, D., Flinn, C. and Wiswall, M. (2016), 'Transfers to households with children and child development', *The Economic Journal* **126**(596), F136–F183.

Del Bono, E., Francesconi, M., Kelly, Y. and Sacker, A. (2016), 'Early maternal time investments and early child outcomes', *The Economic Journal* **126**, F96–F135.

Droogsma, R. A. (2007), 'Redefining hijab: American muslim women's standpoints on veiling', *Journal of Applied Communication Research* **35**(3), 294–319.

Drydakis, N. (2013), 'The effect of ethnic identity on the employment of immigrants', *Review of Economics of the Household* **11**(2), 285–308.

Fan, X. and Wu, L. (2022), The shaping of a gender norm: Marriage, labor, and foot-binding in historical china, Working paper.

Fernández-Reino, M., Di Stasio, V. and Veit, S. (2022), 'Discrimination unveiled: A field experiment on the barriers faced by muslim women in germany, the netherlands, and spain', *European Sociological Review* .

Fouka, V. (2020), 'Backlash: The unintended effects of language prohibition in U.S. schools after World War I', *The Review of Economic Studies* **87**(1), 204–239.

Fryer, R. G. and Levitt, S. D. (2004), 'The causes and consequences of distinctively black names', *The Quarterly Journal of Economics* **119**(3), 767–805.

Galichon, A., Kominers, S. D. and Weber, S. (2019), 'Costly concessions: An empirical framework for matching with imperfectly transferable utility', *Journal of Political Economy* **127**(6), 2875–2925.

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

Garre, F. G., Vermunt, J. K. and Croon, M. A. (2002), 'Likelihood-ratio tests for order-restricted log-linear models: A comparison of asymptotic and bootstrap methods', *Metodologa de las Ciencias del Comportamiento* **4**(1), 1–18.

Gaspard, F. and Khosrokhavar, F. (1995), *Le foulard et la République*, FeniXX.

Gibbons, R. and Waldman, M. (2004), 'Task-specific human capital', *The American Economic Review* **94**(2), 203–207.

Greene, W. H. (2008), *Econometric analysis*, 6th edn, Pearson Education, Inc.

Gresh, A. (2020), 'Islam: A one-sided debate'.
URL: *https://orientxxi.info/dossiers-et-series/islam-a-one-sided-debate,4281*

Guiso, L., Sapienza, P. and Zingales, L. (2006), 'Does culture affect economic outcomes?', *Journal of Economic Perspectives* **20**(2), 23–48.

Gulesci, S., Jindani, S., La Ferrara, E., Smerdon, D., Sulaiman, M. and Young, H. (2021), A stepping stone approach to understanding harmful norms, CEPR Discussion Paper No. DP15776.

Harrison, K. A. (2016), 'Hiding under the veil of" dress policy": Muslim women, hijab, and employment discrimination in the united states', *Geo. J. Gender & L.* **17**, 831.

Heckman, J., Pinto, R. and Savelyev, P. (2013), 'Understanding the mechanisms through which an influential early childhood program boosted adult outcomes', *American Economic Review* **103**(6), 2052–86.

Hu, Z. and Wang, Z. (2021), Nutrition, labor supply, and productivity: Evidence from ramadan in indonesia, Working paper.

Iannaccone, L. R. (1990), 'Religious practice: A human capital approach', *Journal for the Scientific Study of Religion* **29**(3), 297–314.

Iannaccone, L. R. (1992), 'Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives', *Journal of Political Economy* **100**(2), 271–291.

Iannaccone, L. R. (1998), 'Introduction to the economics of religion', *Journal of Economic Literature* **36**(3), 1465–1495.

Institut Français d'Opinion Publique [IFOP] (2019), Etude auprès de la population musulmane en france, 30 ans après l'affaire des foulards de creil, Technical report.

Institut Montaigne (2014–2021), Religion au travail: croire au dialogue. baromètre du fait religieux en entreprise 2014–2021, Technical report.

Institut Montaigne (2016), A french islam is possible, Technical report.

Iyer, S. (2016), 'The new economics of religion', *Journal of Economic Literature* **54**(2), 395–441.

Jia, R. and Persson, T. (2021), 'Choosing ethnicity: The interplay between individual and social motives', *Journal of the European Economic Association* **19**(2), 1203–1248.

Jouili, J. S. (2020), *Pious Practice and Secular Constraints*, Stanford University Press.

Katz, D., Baptista, J., Azen, S. and Pike, M. (1978), 'Obtaining confidence intervals for the risk ratio in cohort studies', *Biometrics* **34**(3), 469–474.

Khosrokhavar, F. (2004), L'islam des jeunes filles en france, *in* 'Le foulard islamique en questions', Éditions Amsterdam, pp. 89–94.

Kirkebøen, L., Leuven, E. and Mogstad, M. (2022), 'College as a marriage market', *NBER Working Paper* (28688).

La Boussole (2004), *Annuaire des mosquées de France*.

Lecomte, J.-P. (2001), Représentations et réalités des fonctions sociales du service militaire dans la société française (1868–2001). Thèse de doctorat en Science politique soutenue en 2001 à Paris, Institut d'études politiques.

Livny, A. (2020), *Trust and the Islamic advantage: religious-based movements in Turkey and the Muslim world*, Cambridge University Press.

Maillard, D. (2017), *Quand la religion s'invite dans l'entreprise*, Fayard.

Maurin, E. and Navarrete-Hernandez, N. (2023), 'Behind the veil: the effect of banning the islamic veil in schools', *Economic Policy* **38**(113), 63–98.

McFadden, D. (1973), Conditional logit analysis of qualitative choice behavior, *in* P. Zarembka, ed., 'Frontiers in Econometrics', Academic Press, New York, chapter 4, pp. 105–142.

Montgomery, J. D. (2010), 'Intergenerational cultural transmission as an evolutionary game', *American Economic Journal: Microeconomics* **2**(4), 115–136.

Morsy, M. (1989), *Les femmes du Prophète*, FeniXX.

Munshi, K. (2011), 'Strength in numbers: Networks as a solution to occupational traps', *The Review of Economic Studies* **78**(3), 1069–1101.

Munshi, K. (2019), 'Caste and the Indian economy', *Journal of Economic Literature* **57**(4), 781–834.

Nordmann, C. (2004), *Le foulard islamique en questions*, Éditions Amsterdam.

Novak, L. (2020), 'Persistent norms and tipping points: The case of female genital cutting', *Journal of Economic Behavior & Organization* **177**, 433–474.

Patacchini, E. and Zenou, Y. (2016), 'Social networks and parental behavior in the intergenerational transmission of religion', *Quantitative Economics* **7**, 969–995.

Patel, D. S. (2012), 'Concealing to reveal: The informational role of islamic dress', *Rationality and Society* **24**(3), 295–323.

Platteau, J.-P. (2017), *Islam instrumentalized*, Cambridge University Press.

Read, J. G. and Bartkowski, J. P. (2000), 'To veil or not to veil? a case study of identity negotiation among muslim women in austin, texas', *Gender & society* **14**(3), 395–417.

Roth, A. E. and Sotomayor, M. A. O. (1990), *Two-sided matching: A study in game-theoretic modeling and analysis*, Cambridge University Press.

Sakalli, S. O. (2019), Secularization and religious backlash: Evidence from turkey, Technical report, Technical Report, Working Paper.

Saleh, M. (2016), 'Public mass modern education, religion, and human capital in twentieth-century Egypt', *The Journal of Economic History* **76**(3), 697–735.

Saleh, M. (2018), 'On the road to heaven: Taxation, conversions, and the Coptic–Muslim socioeconomic gap in Medieval Egypt', *The Journal of Economic History* **78**(2), 394–434.

Scott, J. W. (2009), *The politics of the veil*, Vol. 7, Princeton University Press.

Sénat (2021), 'Compte rendu intégral: Séance du mardi 30 mars 2021', *Journal Officiel de la République française* .

Shofia, N. M. (2020), Why veil? religious headscarves and the public role of women, Working paper.

Silhouette-Dercourt, V., Sy, O. S. and Desjeux, D. (2019), 'Cosmopolitan veiling in paris: Young french muslim women in transition', *Youth and Globalization* **1**(1), 65–87.

Siow, A. (2015), 'Testing Becker's theory of positive assortative matching', *Journal of Labor Economics* **33**(2), 409–441.

Skrondal, A. and Rabe-Hesketh, S. (2009), 'Prediction in multilevel generalized linear models', *Journal of the Royal Statistical Society Series A: Statistics in Society* **172**(3), 659–687.

Squicciarini, M. (2020), 'Devotion and development: Religiosity, education, and economic progress in 19th-century France', *The American Economic Review* **110**(11), 3454–3491.

Valfort, M.-A. (2020), 'Anti-muslim discrimination in france: Evidence from a field experiment', *World Development* **135**, 105022.

van der Hasselt, G. (2019), 'The muslim veil in france: Why so controversial? three questions to hakim el kaouri'.
URL: *https://www.institutmontaigne.org/en/blog/muslim-veil-france-why-so-controversial*

Weber, M. (2013), *The Protestant Ethic and the Spirit of Capitalism*, Routledge.

Weichselbaumer, D. (2020), 'Multiple discrimination against female immigrants wearing headscarves', *ILR Review* **73**(3), 600–627.