

AVERTISSEMENT

Ce document est le fruit d'un long travail approuvé par le jury de soutenance et mis à disposition de l'ensemble de la communauté universitaire élargie.

Il est soumis à la propriété intellectuelle de l'auteur : ceci implique une obligation de citation et de référencement lors de l'utilisation de ce document.

D'autre part, toute contrefaçon, plagiat, reproduction illicite de ce travail expose à des poursuites pénales.

Contact : portail-publi@ut-capitole.fr

LIENS

Code la Propriété Intellectuelle – Articles L. 122-4 et L. 335-1 à L. 335-10

Loi n° 92-597 du 1^{er} juillet 1992, publiée au *Journal Officiel* du 2 juillet 1992

<http://www.cfcopies.com/V2/leg/leg-droi.php>

<http://www.culture.gouv.fr/culture/infos-pratiques/droits/protection.htm>



THÈSE



En vue de l'obtention du

DOCTORAT DE L'UNIVERSITE DE TOULOUSE

Délivré par l'Université Toulouse Capitole

École doctorale : **Sciences Economiques-Toulouse School of Economics**

Présentée et soutenue par

RAIBER Eva

le 18 juin 2019

Essays in Applied Microeconomics

Discipline : **Sciences Economiques**

Unité de recherche : **TSE-R (UMR CNRS 5314 – INRA 1415)**

Directeur de thèse : Monsieur Paul SEABRIGHT, Professeur,
Toulouse School of Economics (IAST)

JURY

Rapporteurs Madame Eliana LA FERRARA, Professeure, Università Bocconi
Madame Sylvie LAMBERT, DR INRA et Professeure, Paris School of
Economics

Suffragants Madame Emmanuelle AURIOL, Professeure, Toulouse School of Economics
(IAST)

Declaration

I gratefully acknowledge financial funding for the first three years of my Ph.D. from the Université de Toulouse 1 Capitole and the two last years from the ANR - Labex IAST.

Acknowledgement

I am foremost grateful to my thesis advisor, Paul Seabright, who from the beginning believed in my ideas and believed in me. I feel very privileged for all the time and resources he invested in our projects, my ideas and all the additional opportunities along the way. He taught me to go for the interesting questions, be confident about what I am doing, but also listen to others and recognize when I make mistakes. I am extremely thankful to Emmanuelle Auriol for always being approachable, supportive and positive. I believe that every graduate student should have someone who cheers them on. I am indebted to Chris Udry who welcomed me at Northwestern, let me be part of his amazing team and was very generous with his time and advice. I am also thankful to Eliana La Ferrara and Sylvie Lambert for agreeing to be on my jury, for their time and comments.

I am furthermore grateful to my great co-authors Jeanne Bovet, Weiwei Ren, Charlotte Wang, Amma Panin and Julie Lassébie. My work has immensely benefited from valuable discussions with Silvain Chabé-Ferret, Astrid Hopfensitz, Francois Poinas, Thierry Magnac, Matthias Doepke, Stéphane Straub, Nancy Qian, Abigail Adams, Sylvie Démourger, Ingela Alger, Marie-Clarie Villéval, Neil Ferguson and countless others from the Toulouse School of Economics and at the many different conferences and seminars. I am thankful to Tilman Brück, Dorothea Kübler and Rajshri Jayaraman for welcoming me at their respective research group. I am also indebted to Wilfried Sand-Zantman and the staff at the Toulouse School of Economics and the Institute for Advanced Study in Toulouse for their support. A special thank you goes to all the research assistants and enumerators that I worked with, as well as to all the survey and experiment participants, that are the true foundation of this thesis.

My time in Toulouse and partly in Berlin and at Northwestern would not have been the same without my friends and colleagues. I am extremely grateful to have met so many great people on this way and I am looking forward to seeing where life takes us all. I am thankful to my family who was always there when I needed them with a special thanks to Klara Raiber who is as much a colleague as a family member. My final thank you goes to Hannes Oehme who believed in me, supported me and endured all the ups and downs along the way.

Abstract

This thesis consists of three chapters in applied microeconomics. Two chapters fall into the field of family economics and one into the field of the economics of religion.

The first chapter asks if anticipated future fertility affects educational investment. Theory suggests that the number of children planned in the future can affect the returns to education, the resources available for family consumption and the incentives to find a partner. This chapter uses varying eligibility criteria for second child permits during the One-Child Policy in China as a natural experiment, which provides plausible exogenous variation in the cost of the second child. I use second child permits that are conditional on time-invariant individual characteristics and show that they have a strong positive effect on the likelihood of having a second child between 1990 and 2005. They are therefore expected to change anticipated fertility among compliers. I find that fulfilling an eligibility criterion at secondary school age increases the time invested in education and the likelihood of continuing schooling after middle school. The effect appears concentrated in the subset of compliers: individuals who increase their anticipated number of children as a response to eligibility. It can be explained by the high cost of raising children, by the second child having no or only a short-term effect on parental labour supply and by concerns about finding a spouse.

The second chapter, co-authored with Weiwei Ren, Jeanne Bovet, Paul Seabright and Charlotte Wang, investigates marriage patterns and underlying preferences in China. We estimate mate preferences based on the evaluation of a series of randomly created profiles and connect our results to recent marriage patterns in the general population. Data is collected on parents (or other relatives) who search for a spouse on behalf of their unmarried adult child and on unmarried students. We confirm that male profiles with a high income and real estate ownership have a higher probability of being selected by parents than profiles with low income and no real estate, but not female profiles. We find that parents dislike profiles with less education than their son or daughter, but we do not find evidence for a dislike of female profiles with high education. However, some parents may have excessively high expectations about the educational level of their son-in-law. Parents mostly prefer male profiles with a similar age to their daughter and female profiles that are younger than their son. Yet, when their son is in his twenties, parents also accept women who are slightly older. If marriages formed according to parents' preferences, the most common case would be the husband being 2 or 3 years older. Yet, in the general population, the most common case for recent marriages

is spouses having the same age. This would be the case if marriages formed according to students' preferences. Both parents' and students' preferences predict a high degree of assortativeness on education, which we observe in recently married couples in the overall population.

The third chapter, co-authored with Emmanuelle Auriol, Julie Lassébie, Amma Panin and Paul Seabright, provides experimental support for the hypothesis that insurance can be a motive for religious donations by members of a Pentecostal church in Ghana. We randomize enrolment into a commercial funeral insurance policy, then church members allocate money between themselves and a set of religious goods in a series of dictator games with significant stakes. Members enrolled in insurance give significantly less money to their own church compared to a control group of members that only receive information about the insurance. Enrolment also reduces giving towards other spiritual goods. We set up a model exploring different channels of religiously based insurance. The implications of the model and the results from the dictator games suggest that adherents perceive the church as a source of insurance and that this insurance is derived from beliefs in an interventionist God. Survey results suggest that material insurance from the church community is also important and we hypothesize that these two insurance channels exist in parallel.

Résumé

Cette dissertation est composée de trois chapitres en microéconomie appliquée. Deux chapitres s'inscrivent dans le champ de l'économie de la famille et un chapitre dans le champ de l'économie de la religion.

Le premier chapitre pose la question de savoir si la fertilité future anticipée a un effet sur l'investissement dans l'éducation. Un modèle théorique suggère que le nombre des enfants planifié dans le futur peut affecter les rendements de l'éducation, les ressources disponibles pour la consommation familiale et la motivation pour chercher un conjoint. Ce chapitre utilise l'expérience naturelle constituée par les variations dans les critères d'éligibilité pour le permis d'avoir un deuxième enfant pendant la politique d'enfant unique en Chine, laquelle fournit une variation plausiblement exogène dans le coût d'avoir un deuxième enfant. J'utilise les permis d'avoir un deuxième enfant conditionnés par des caractéristiques individuelles invariables dans le temps et je montre que ces permis ont un fort effet positif sur la probabilité d'avoir un deuxième enfant entre 1990 et 2005. Ils sont donc susceptibles de changer la fertilité anticipée parmi les *compliers* (conformistes). J'observe que satisfaire à un critère d'éligibilité à l'âge de l'école secondaire augmente le temps investi en éducation et la probabilité de poursuivre les études après l'éducation intermédiaire. L'effet paraît concentré dans le groupe des *compliers* - les individus qui augmentent leur nombre d'enfants anticipé en réponse à leur éligibilité. Ce résultat peut être expliqué par un coût élevé d'éducation des enfants, par le fait que le deuxième enfant n'ait qu'un effet de court terme sur l'offre de travail parental, et par le souci de trouver une conjointe.

Le deuxième chapitre, coécrit avec Weiwei Ren, Jeanne Bovet, Paul Seabright et Charlotte Wang, examine les schémas matrimoniaux et les préférences sous-jacentes en Chine. Nous estimons les préférences conjugales par l'évaluation d'une série de profils créés aléatoirement et nous relierons les résultats aux schémas matrimoniaux récents dans la population générale. Les données sont recueillies auprès des parents (ou d'autres membres de famille) qui cherchent un ou une partenaire pour leur enfant adulte non marié, ainsi qu'auprès des étudiants non mariés. Nous confirmons que les profils masculins avec un haut revenu et une propriété immobilière ont une probabilité plus élevée d'être sélectionnés par les parents que les profils avec un faible revenu et sans propriété immobilière. Mais il n'en va pas de même pour les profils féminins. Nous constatons que les parents n'aiment pas les profils avec un niveau d'éducation inférieure que celui de leur fils ou leur fille, mais

nous ne trouvons aucune preuve que les parents seraient défavorables à des profils féminins avec un niveau d'éducation plus élevé. En revanche, certains parents peuvent avoir des attentes excessivement exigeantes quant au niveau d'éducation de leur beau-fils. Les parents préfèrent généralement les profils masculins avec un âge similaire à celui de leur fille et les profils féminins qui sont plus jeunes que leur fils. Toutefois, lorsque leur fils a la vingtaine, les parents acceptent aussi des femmes légèrement plus âgées. Si les mariages se formaient selon les préférences des parents, le cas le plus courant serait celui d'un mari avec deux ou trois ans de plus que sa femme. Cependant, dans la population générale, le cas le plus courant parmi les mariages récents est celui des conjoints du même âge. Ce serait le cas si les mariages se forment selon les préférences des étudiants. Les préférences des parents et celles des étudiants prédisent un degré élevé d'appariement assortatif selon le niveau d'éducation, ce que nous observons chez les couples récemment mariés dans la population générale.

Le troisième chapitre, coécrit avec Emmanuelle Auriol, Julie Lassébie, Amma Panin et Paul Seabright, étudie dans un cadre expérimental l'hypothèse selon laquelle les dons religieux des membres d'une église pentecôtiste au Ghana auraient un rôle d'assurance contre les chocs négatifs. Nous offrons aléatoirement à certains membres de la communauté religieuse une souscription à une assurance funéraire disponible sur le marché, et nous mesurons en laboratoire comment les individus choisissent de répartir une somme d'argent entre eux-mêmes et un ensemble de bénéficiaires religieux et spirituels. Les membres ayant reçu l'assurance donnent significativement moins d'argent à leur propre église par rapport au groupe de contrôle des membres qui ont seulement reçu des informations à propos de l'assurance. Les premiers réduisent également leurs dons aux autres bénéficiaires spirituels. Nous développons un modèle qui explore les différents mécanismes d'assurance fondés sur les dons religieux. Les résultats théoriques et expérimentaux suggèrent que les participants à notre étude perçoivent l'église comme une source d'assurance et que cette assurance est dérivée de croyances dans un Dieu interventionniste. Les participants à l'étude déclarent également qu'une assurance matérielle émanant de la communauté est également importante et nous émettons l'hypothèse que ces deux mécanismes d'assurance existent en parallèle.

Contents

Declaration	i
Acknowledgement	ii
Abstract	iii
Résumé	v
Introduction	1
1 Anticipated Fertility and Educational Investment: Evidence from the One-Child Policy in China	4
1.1 Introduction	6
1.1.1 Links with the Literature	8
1.2 Context: The One-Child Policy, second child permits and education in China	9
1.3 Model	12
1.3.1 Set-up	13
1.3.2 Optimal educational level and optimal number of children	15
1.3.3 Effect of 2nd child permits	15
1.3.4 Gender differences and marriage	19
1.3.5 The effect of fines on education	21
1.4 Data	21
1.4.1 Individual Data	21
1.4.2 Policy exemptions	22
1.5 Empirical methodology	24
1.5.1 Effect of eligibility on education	24
1.5.2 Effect of eligibility on fertility outcome	25
1.5.3 Isolating the effect on <i>increasers</i>	26
1.6 Results	27
1.6.1 Effect of second child permits on education	27
1.6.2 Effect of eligibility on fertility outcome	33
1.6.3 Isolating the effect on <i>increasers</i>	34

1.6.4	Effect of eligibility and the sex ratio	37
1.7	Discussion: Why would those that expect more children educate more?	37
1.7.1	Cost of raising children	38
1.7.2	Old-age support and the double burden of parents	39
1.7.3	Parental labour supply and earnings	40
1.8	Conclusion	40
1.9	Appendix	43
1.9.1	Figures	43
1.9.2	Tables	50
1.9.3	Model appendix	59
1.9.4	Further empirical results	62
2	Spousal Preferences and Parental Search: Data from China	67
2.1	Introduction	69
2.1.1	Links with the literature	72
2.2	Context	73
2.2.1	Marriage in China	73
2.2.2	Marriage patterns in contemporary China	75
2.3	Parental Search at Green Lake Park	77
2.3.1	Who is searching?	78
2.3.2	How are they searching?	79
2.3.3	Why are they searching?	79
2.3.4	On whose behalf are they searching?	80
2.3.5	Comparison: The student sample	81
2.4	Preference Estimation	82
2.4.1	Hypothetical profiles	82
2.4.2	Estimation Framework	83
2.4.3	Results: Parents	84
2.4.4	Results: Students	89
2.4.5	What do we learn from the preferences?	89
2.5	Comparison with marriage patterns	92
2.5.1	Stated acceptance limits	92
2.5.2	Match Simulation	98
2.5.3	Discussion	100
2.6	Conclusion	105
2.7	Appendix	106
2.7.1	Figures	106
2.7.2	Tables	109

3	God insures those who pay? Formal insurance and religious offerings in Ghana	115
3.1	Introduction	117
3.1.1	Links with the literature	119
3.2	Context	121
3.2.1	Pentecostalism	121
3.2.2	The role of the church in Pentecostalism	123
3.2.3	Giving to God to access insurance	124
3.3	Experiment	125
3.3.1	Enrolment treatment: Funeral insurance policy	125
3.3.2	Control group and no insurance information treatment	126
3.3.3	Recruitment	127
3.3.4	Experimental Setup and Design	127
3.4	Model	129
3.4.1	Setup	130
3.4.2	Optimal giving to the church in the absence of any insurance	130
3.4.3	Community insurance: optimal giving when giving reduces the size of a loss	131
3.4.4	Spiritual insurance: optimal giving when giving reduces the subjective probability of a loss	132
3.4.5	Experimental hypotheses	133
3.5	Experimental Results	134
3.5.1	Descriptive statistics	134
3.5.2	Summary of the allocation decisions	135
3.5.3	Treatment effects	137
3.5.4	Discussion	142
3.6	Conclusion	144
3.7	Appendix	146
3.7.1	Tables	146
3.7.2	Model Appendix	156
3.7.3	Questionnaire	164
3.7.4	Experimental protocol	173

List of Figures

1.1	Fertility rate in China, 1960 - 2014	10
-----	--	----

1.2	Share of population qualifying for 2nd child permit and incidences of second child permit reforms.	23
1.3	Effect of anticipated eligibility on finishing junior high school, enrollment into senior high school and enrollment into college or university	31
1.4	The effect of fertility preferences on the number of children.	44
1.5	Effect size of anticipated eligibility on likelihood of senior high school enrolment at different age thresholds.	44
1.6	Average sex ratio for those who turned 16 between 1990 and 2000.	45
1.7	Importance of predictors in random forrest estimation	45
1.8	Importance of predictors including the cohort sex ratio (from 2000 census) on the top and monetary fines on the bottom	46
1.9	Total educational expenditure as share of family expenditure	47
1.10	Different items of educational expenditure as share of family expenditure.	47
1.11	Expenditure on education for families with two children and a matched control group with one child	48
1.12	Log monthly working hours in 2010 according to the year of birth of the second child	48
1.13	Log monthly income in 2010 according to the year of birth of the second child	49
2.1	How individuals met their spouse.	74
2.2	Educational difference between husband and wife	75
2.3	Age difference between husband and wife	76
2.4	Age difference and correlation between husband and wife's ages over time	77
2.5	English example for the hypothetical profiles that were shown to respondents.	83
2.6	Female subjects: The predicted likelihood of choosing the profile according to the age difference (age of profile minus age of eligible subject).	87
2.7	Male subjects: The predicted likelihood of choosing the profile according to the age difference (age of profile minus age of eligible subject).	88
2.8	Education: Acceptance Sets	94
2.9	Education: Acceptance Sets with frequency of recent marriages	95
2.10	Acceptance Sets on Age	96
2.11	Acceptance Sets on Age with frequency of recent marriages	97
2.12	Education: Comparing actual educational differences between husband and wife and educational differences in simulated matches.	101
2.13	Age: Comparing actual educational differences between husband and wife and educational differences in simulated matches.	102
2.14	Education: women-proposing Gale-Shapley algorithm	107
2.15	Age: women-proposing Gale-Shapley mechanism	108
3.1	Distribution of giving among normal church population.	136
3.2	Community insurance - Numerical example	159

3.3	Community insurance - Threshold on θ	160
3.4	Spiritual insurance - Numerical example	163
3.5	Spiritual insurance: Response of optimal giving to an increase in D	163

List of Tables

1.1	Effect of eligibility at age 16 on the years of education.	29
1.2	Sex difference in the effect of eligibility at age 16 on the years of education.	30
1.3	Effect of eligibility status at age 27/28 on the likelihood of having a second child.	34
1.4	Decomposing the effect of eligibility	36
1.5	Example of when provinces formalized eligibility criteria.	50
1.6	Summary statistics	51
1.7	The effect of the eligibility on the likelihood of completing a degree.	52
1.8	Placebo test.	53
1.9	Effect of anticipated eligibility on the years of education - Robustness Checks.	54
1.10	Spouse dependence in the eligibility criteria.	55
1.11	Sex ratio and eligibility.	56
1.12	Effect of the second child policy on educational expenditure.	57
1.13	Effect of a second child on parental working hours.	57
1.14	Effect of the second child policy on household balance.	58
1.15	Monetary fines and eligibility.	64
1.16	Effect of eligibility on age at marriage and age at first birth.	66
2.1	Comparison between STEP and QSAMPY data	81
2.2	Estimation using hypothetical choice data and the question which profile the respondent would like their eligible subject to meet.	85
2.3	Estimation using hypothetical choice data and the question which profile the respondent would like their eligible subject to meet.	86
2.4	Student Sample: Estimation using hypothetical choice data and the question which profile the respondent would like to meet.	90
2.5	Comparing real marriage distribution with predicted distributions	103
2.6	Educational Levels of husband and wife	109
2.7	Summary statistics on respondents	109
2.8	Share of respondent indicating reason for search as important or very important.	110
2.9	Summary statistics on eligible subjects on whose behalf respondents are searching	110

2.10	Estimation using hypothetical choice data and the question which of the two profiles the respondent prefers for the eligible subject.	111
2.11	Lower and upper acceptance bounds on education according to parents	112
2.12	Lower and upper acceptance bounds on age according to parents	113
2.13	Estimations to simulate marriage outcomes	114
3.1	Pairs of dictator game recipients.	128
3.2	Giving after receiving insurance information compared to giving with no insurance nor insurance information	138
3.3	Giving after enrolment in insurance compared to giving with insurance information only	140
3.4	Summary statistics of church branches	146
3.5	Social role of the churches	147
3.6	Summary statistics of study participants and comparison with general population . .	147
3.7	Summary statistics of study participants across treatments	148
3.8	Financial role of the churches	149
3.9	Giving after enrolment in insurance compared to giving with no insurance nor insurance information	150
3.10	Comparison of regular and revival week participants	151
3.11	Treatment effects on church giving for total sample (column 1) and revival week only (column 2)	152
3.12	Information on church branches	153
3.13	Non-anonymous church giving	153
3.14	Treatment effects for other dictator games	154
3.15	Comparing revival church members and church members giving relatively high amounts to the church	155

Introduction

Before I started the graduate program in Toulouse, I did an internship at a small public policy company in India. At one point, my colleague explained to me that once she felt ready to get married, her parents would preselect potential suitors. She would then meet him and his family and make the decision together with her family. Marriage is one of the most decisive steps for many women in the world: it is what mostly determines her socioeconomic, psychological and physical well-being. For such an important decision, why would anyone not take into account the advice of their parents who have marriage experience?

Not every parent-child discussion about marriage is as harmonious and collaborative. In many marriages in South Asia, the bride and the groom meet at their wedding for their first time. Another example is my grandmother who had to wait until the age of 21 to get married because her father would not give her his consent. I asked myself, when do parents and children cooperate and when is there a conflict? If there is a conflict, is it over which spousal characteristics are important? In any case, I became interested in the economics of marriage.

While on the search for a topic for my job market paper, my thesis advisor urged me to look at marriages in China. As I describe in section 2.2.1, China used to have a tradition of arranged marriages which was outlawed when the Communist Party came to power. Yet, this does not mean that parents are not involved in the search for a marriage partner. On the contrary, parental involvement might become more important again with the increase in financial inequality: a “good match” might become again more important and worth serious search efforts.

It is impossible to study marriages and families in China without stumbling across the One-Child Policy. This policy is certainly the most invasive large-scale fertility policy to date. As I describe in section 1.1.1, it has not only affected fertility outcomes but also marriage outcomes and ethnic identity. While thinking through how the One-Child Policy would affect marriages, I realized that it could also affect premarital investment, in particular, investment in education. The more general question is if the number of children someone plans to have in the future affects educational investment decision. Chapter 1 addresses this question.

When motivating the relevance of this question, I often think about my parents. My parents met when they were quite young and both wanted a big family. The strategy of my father, at least partially to this effect, was to get as much education as possible. So he got a PhD. The strategy of my mother was to get a university education but in an area where she could exit and enter the labour market relatively easily and where there was a flat earning profile. In chapter 1 I find that in China, men who are allowed to have two children receive more education than men who are allowed to have only one child. The effect on women is less clear, possibly due to heterogeneity in how women’s labour supply changes with the birth of a child.

The original project for which I went to China worked out thanks to the amazing work of my coauthors. We were able to interview parents who are actively looking for a potential spouse for their adult child at a designated area in a public park during weekends. Parents with unmarried children meet, chat and exchange information. As a comparison sample, we interviewed unmarried students at the local university. Our findings are described in chapter 2. Part of the project is published in [Bovet et al. \(2018\)](#). Generally, we find that parents see themselves as agents for their children, however, their preferences do not seem to perfectly overlap. The parent-child dynamics, in particular when they do not fully agree on the preferred characteristics, are an interesting avenue of future research.

The third chapter at first seems a bit at odds with the first two chapters. It was the first project I was involved in when I started my PhD and falls within the field of the economics of religion. Yet, it has very important connections with the other two chapters. While the family is the most important network in many countries, for many, religious networks are a close second. For the empirical analysis in chapter 3, we interviewed members of a Pentecostal church. When asked about who they would ask for counsel if they had personal or family issues nearly 70% indicated that they would ask their pastor and nearly 40% would turn to a family member (multiple answers were possible). When they would need financial help, 51% would turn to their family and 24% to their church community.

The starting point for chapter 3 was the question why religious organisations play such an important role in people lives, in particular in developing countries in Sub-Saharan Africa, where many give between 10 to 15% of their available income to them. We focus on two services that religious organisations potentially provide: financial and “spiritual” insurance, the latter being derived from a belief in a higher power. Yet, the network function of religion is undeniable for other services as well. In our survey, 40% of respondents state that they have found or believe they are most likely to find their spouse in the church community. Nearly 40% say that they do business with other members of the church. While teaching about God is the most important reason they state why they are a member of this particular church, respondents also indicate moral guidance (53%), the atmosphere of the services (38%) and the welcoming community (21%).

In the Pentecostal church that we worked with individuals that have the same concern would meet and pray together. As their children’s marriage is apparently a major concern for many parents, particularly mothers, they would meet for specific prayer groups. Here they are able to exchange and this is essentially what Chinese parents are doing in public parks.

Chapter 1

Anticipated Fertility and Educational Investment: Evidence from the One-Child Policy in China¹

EVA RAIBER

¹This chapter is a revised and extended version of the TSE Working paper “Expected Fertility and Educational Investment: Evidence from the One-Child Policy in China” (Raiber (2017)). I thank Paul Seabright, Chris Udry, Silvain Chabé-Ferret, Mathias Reynaert, Sylvie Démurger, Thierry Magnac, Matteo Bobba, Tilman Brück, Nancy Qian, Matthias Doepke, Paula Gobbi, Florence Nimoh, participants and discussants at the Ruhr Graduate School Doctoral Conference 2017, the ENTER Jamboree at University College of London 2017, the EUDN PhD workshop 2017, the 2017 Lindau Nobel Laureates Meeting, the CSAE 2018, seminar participants at ISDC Berlin, Toulouse School of Economics, ECARES Brussels, Northwestern University and the University of Warwick for valuable feedback and comments. I am grateful to Wanying Zhao for excellent research assistance.

ABSTRACT

Does anticipated future fertility affect educational investment? Theory suggests that the number of children planned in the future can affect the returns to education, the resources available for family consumption and the incentives to find a partner. This chapter uses varying eligibility criteria for second child permits during the One-Child Policy in China as a natural experiment, which provides plausible exogenous variation in the cost of the second child. I use second child permits that are conditional on time-invariant individual characteristics and show that they have a strong positive effect on the likelihood of having a second child between 1990 and 2005. They are therefore expected to change anticipated fertility among compliers. I find that fulfilling an eligibility criterion at secondary school age increases the time invested in education and the likelihood of continuing schooling after middle school. The effect appears concentrated in the subset of compliers: individuals who increase their anticipated number of children as a response to eligibility. It can be explained by the high cost of raising children, by the second child having at most a short-term effect on parental labour supply and by concerns about finding a spouse.

1.1 Introduction

Educational and fertility choices are major life decisions that are deeply interconnected. The microeconomics literature has largely focused on the effect of parental education on fertility (e.g. McCrary and Royer (2011), Osili and Long (2008), Duflo et al. (2015), Lavy and Zablotsky (2015)). The reverse question has rarely been addressed: Do individuals who anticipate having more children receive more or less educational investment than those who anticipate none or fewer? Answering this question is challenging because the number of children one plans to have in the future is usually unobserved. Even if observed, it is correlated with other variables that affect the demand for education.

Faced with ageing societies, policy-makers in many high-income countries are eager to encourage childbearing. Child benefits, free childcare and paid parental leave are policies implemented or discussed in most low fertility countries. They incentivize childbirth by reducing the costs associated with having children. How do such reforms influence educational investment? Do they have the side effect of keeping future parents from investing in education? Theory suggests that children can affect parental labour supply and the income available for family consumption. The desire to have children can also strengthen the incentives to find a partner. Therefore, the anticipated number of children can have an effect on the returns to education and the marginal utility of income. The sign of the overall effect is ambiguous and needs to be evaluated empirically.

This chapter uses variation in the strictness of the One-Child Policy in China to estimate the effect of anticipated fertility on educational investment. While one child per family was the norm, having another child was allowed with a “second child permit”. The eligibility criteria for these permits were set on the provincial level. Policy changes provide plausible exogenous variation in the costs of having a second child. Several criteria were conditional on observable characteristics that are time-invariant. These include ethnicity, household registration status and being an only child. It allowed their families to predict if their teenage child would be eligible in the future. If second child permits have a strong effect on fertility outcomes, they are expected to have an effect on anticipated fertility.

The criteria for second child permits varied over time and between provinces. They are not directly related with provincial educational policies. This allows the use of difference-in-differences (DID) and double DID identification strategies (also called triple differences approach). I calculate whether individuals fulfil a second child permit criterion at the age when they usually finish mandatory secondary school. At this point, they and probably mostly their family must decide whether to continue with voluntary senior high school. I find that fulfilling a second child permit criterion significantly increases schooling. Between 1990 and 2005, eligible individuals spend on average around 0.9 years more on education. The likelihood of enrolment into higher secondary

school is around 9 percentage points higher. The effect is stronger for men (but not significantly so); the effect for women is positive but insignificant at around 0.6 years of education.

To affect anticipated fertility, permits need to have an actual impact on fertility outcomes. I use the same second child permit criteria as before and calculate the eligibility status at the mean age when individuals decide to have a second child, for an older sample. I find that eligibility increases the likelihood of having a second child by around 14 percentage points. This is an important effect, yet it shows that not all individuals change their fertility decision with eligibility.

Individuals and their families might plan two children independent of their eligibility status. Yet, they still benefit from not having to pay monetary fines once they are eligible. I use those not eligible in the older sample as the training sample for a random forest estimation. It predicts if individuals would have a second child even when they are not eligible. As expected, the fertility outcome of those predicted to have a second child without permits is not influenced by eligibility. I also find that for them the effect of eligibility on enrolment into higher secondary school is significantly lower. This suggests that the positive effect on schooling comes from compliers: those who increase their anticipated fertility from one child to two as a response to becoming eligible. Further treatment heterogeneity analysis shows that men in provinces with a skewed sex ratio are the most affected.

Primary and secondary schooling decisions are typically made before individuals have children. However, individuals, together with their own parents, might take into account how many children they plan to have in the future. In the context of secondary school decisions, parents are probably the main decision-makers. They might also be more forward-looking and aware of the consequences of having children. First, children can be an important cost factor that decreases the monetary resources available for family consumption. To prepare for this cost in the future, families can use educational investment to smooth consumption over time (intertemporal consumption smoothing effect). This channel is particularly important when individuals have to financially support their own parents at the same time as raising their children. Second, children can also affect the time their parents spend working in the labour market. The income loss increases in the parents' educational level and has an effect on lifetime returns to education. The less time the parent can spend in the labour market, the lower are the pay-offs from education (labour supply effect). Third, individuals who want to have children in the future might have a higher incentive to find a partner. If the sex ratio is skewed, individuals of the more frequent sex compete for partners. They can use education to increase their chances (marriage market effect). In section 1.3, I use a two-period model to demonstrate these channels and highlight the assumptions needed.

The empirical findings suggest that higher anticipated fertility can increase educational investment. Thus, policies that encourage childbearing can have positive side effects on education.

However, labour market conditions are important. I show that if there are not labour supply consequences of children, higher costs of raising children increases the effect of anticipated fertility on education. Wanting children does not have a negative effect on education if parents can quickly return to the labour market after childbirth. It is important that they do not have to fear lower returns to education than their childless co-workers. I show that the effect of anticipated fertility on education is decreasing in the loss of working time due to childcare.²

Finally, I provide descriptive statistics about the cost of children and the parental labour supply in China. Nowadays, raising children is considered very costly. For a family with two children, educational expenditure accounts for up to 30% of available family income. Yet, parental labour supply and income are not affected in the long-run by having a second child. Women with a second child work and earn less than women with only one child for at most 3 years and then catch up. Other studies find no or little evidence on the negative effects of fertility on parental labour supply (Guo et al. (2017), He and Zhu (2016)).

1.1.1 Links with the Literature

Theoretical growth models and country level empirical work usually connect low fertility rates and high human capital investment (Becker et al. (1990), Rosenzweig (1990), Kalemli-Ozcan (2003)). Within a given setting, having many children is often correlated with lower levels of parental education. Several papers focus on identifying the causal effect of education on fertility outcome. They usually use shocks to the supply of education to address the issue of unobserved variables being correlated with both fertility and education and the decisions about them being taken simultaneously (e.g. Duflo et al. (2017), Osili and Long (2008), Fort et al. (2016)). In this chapter, I use a shock to the cost of children to look at the reverse question.

Female education is often associated with lower fertility rates (Osili and Long (2008), Lam and Duryea (1999), Schultz (1997), Duflo et al. (2015), Duflo et al. (2017), Lavy and Zablotsky (2015)). The main economic argument is that the opportunity costs of having a child for an educated woman are higher than for a non-educated woman (Becker (1981)). Education can also increase the knowledge of contraception methods (Rosenzweig and Schultz (1989)) and the bargaining power of women who want fewer children than men (Manser and Brown (1980)). By raising women's income, education can shift the fertility choice toward fewer children of higher quality (Becker (1960), Becker and Lewis (1973), Willis (1973)). This result, however, does not necessarily hold for developed countries (Fort et al. (2016), McCrary and Royer (2011), Monstad et al. (2008)).

An important share of papers in the fertility and education literature focus on the quality-quantity trade-off. In the theoretical framework based on Becker and Lewis (1973), parents trade-off between the number of children and how much they invest in each child. The One-Child Policy

²This is shown under constant absolute risk aversion.

in China has been used to evaluate this trade-off empirically (Qin et al. (2016), Li and Zhang (2016), Rosenzweig and Zhang (2009), Qian (2009)). This chapter goes one step further. Given the number of children, parents take into account the number of grand-children when they make investment decisions.³

Recent work by Huang et al. (2016b) is the most closely related. However, they only use regional variation in monetary fines that I find has no effect on either schooling investment or on fertility choice. They also assume that an increase in monetary fines decreases anticipated fertility and do not discuss that it might increase the cost of a second child while keeping anticipated fertility constant.


The exemptions for ethnic minorities have been used to study inter-ethnic marriages (Huang and Zhou (2015)) and ethnic identity (Jia and Persson (2017)). This implies that on top of having fertility consequences, the One-Child Policy shaped many other socioeconomic decisions. This chapter strengthens this conclusion.

1.2 Context: The One-Child Policy, second child permits and education in China

Empirical identification relies on exogenous variation in the cost of having another child. The One-Child Policy (OCP) in China between 1979 and 2015 was based on the goal of one child per family, setting out fines and penalties for the birth of a second child. At the same time, provincial governments could issue permits for the second child and in rare cases for the third. The criteria under which one could apply for such a permit changed over time. This section describes the policy and its regulations, as well as the functioning, motivation and implications of second child permits.

Family planning has been of particular importance to the Chinese government for the past decades. After the Great Famine (1959-1961), the central government promoted ambitious family planning policies. These include the “Later, Longer, Fewer” campaign from 1971-1979, the OCP from 1979-2015 and the recent Two-children policy. During the “Later, Longer, Fewer” campaign, the government promoted later marriage, longer birth intervals and fewer children. One child per family was optimal and two was acceptable for urban couples, three for rural couples. Penalties were introduced for those who did not comply (Whyte et al. (2015)). Birth control and abortion were promoted. The campaign also included a strong element of coercion. Women, mainly in rural areas, were pressured to abort out-of-quota children and to get sterilized after the birth of the third

³I assume that individuals can plan their fertility outcome and the timing of their pregnancies. This is an appropriate assumption for many high and middle-income countries, including China. A distinct but connected strand of literature looks fertility uncertainty. Contraceptive methods give women certainty over the pregnancy consequences of sex and thus decrease the risk of tertiary schooling investment (Goldin and Katz (2002), Ananat and Hungerman (2012), Miller (2010)).



Source: World Bank. It uses the World Bank definition of Total Fertility Rate. It is calculated as the average number of children that would be born to a woman if she were to experience the exact current age-specific fertility rates through her life.

Figure 1.1: Fertility rate in China, 1960 - 2014

child (Whyte et al. (2015)). Fertility rates in China fell sharply during this period (see figure 1.1), though there is no consensus about how much of the fall was due to measures specifically targeting fertility. The OCP was introduced after the sharp drop in fertility rates.

Between 1978 and 1980, the central government introduced the goal of one child per family. Provinces were to implement this goal by setting fines for the birth of a second child and by providing birth control measures. Between 1979, the policy was rolled out on a county-to-county basis (Qian (2009), Almond et al. (2019)). However, in rural areas where the one-child limit met significant resistance, implementation was delayed (Baochang et al. (2007)). In 1983, the OCP precipitated a huge large wave of abortions and sterilizations (Whyte et al. (2015))⁴.

Between 1982 and 1984, provincial governments started to issue more or less formal guidelines under which conditions married or remarried couples could apply for a second child permit. Thereby, they relaxed the one child per family limit significantly (Scharping (2013)). Between 1986 and 1991, provincial governments produced official family planning regulations which outlaid in detail the different criteria. Most of them were revised at least once in the 1990s and again after 2001. While transparency about family policies increased, the use of coercive measures such as forced abortion and sterilization dropped substantially (Whyte et al. (2015)).

The 1980s represent the time of introduction and adjustment of the policy, whereas the 1990s and 2000s represent a time of stability from a family policy perspective. Transparency of the policy measures is important to form adequate fertility expectations. This motivates using the time frame of 1990 to 2005 for the empirical evaluation. Furthermore, the 1980s coincide with the implementa-

⁴14.4 million abortions, 20.7 million sterilizations and 17.8 million IUD insertions (Whyte et al. (2015))

tion of the Law on Nine-Year Compulsory Education. The law was introduced to attain universal education for nine years: six years at primary school (from age 6/7 to 12/13) and three years at junior high school (from age 12/13 to 15/16). After junior high school, students can voluntarily continue with general or vocational senior high school. Therefore, the first and most relevant educational investment decision is taken when the child is around 16 years old.

During the OCP, couples that wanted to have a child had to apply for a permit allowing them to do so and only married couples were able to apply. Couples that had a second child without a second child permit officially had to pay monetary fines. These were set as a function of the couple's income (Scharping (2013)). Couples with higher income thus had to pay higher fines in absolute terms. Additionally, parents potentially faced non-monetary penalties. These included losing their job or having their career opportunities restricted. There is no accessible data on the enforcement of monetary fine and how frequent other social penalties were. Couples officially had to obtain the second child permit before having the second child. However, this posed a significant financial burden to local rural governments. Second child permits were presumably given out after the birth if the couple fell into an exemption category (Scharping (2013)).

Eligibility criteria for second child permits varied on the provincial level and between rural and urban areas. The household registration status (*hukou*) was an important determinant for eligibility. It is either agriculture/rural or non-agricultural/urban. Couples could only apply for a child permit at their place of household registration. This restricted strategic between-province and urban-rural migration. Most exemptions also required that the applicant respect late childbirth (birth of the first child after age 24 for women) and an acceptable birth interval (between 4 and 7 years).

Provinces introduced several different exemptions over time for different reasons. The most known is the policy that allowed couples in rural areas whose first-born was a girl to have a second child. In five provinces, couples living in rural areas were always allowed to have two children (Baochang et al. (2007))⁵. These exemptions were introduced to appease the rural population and to ensure sufficient labour in the agricultural sector.

Couples from ethnic minorities were often allowed to have two children or were even completely exempted from the policy. However, this depended on the province, if the couple lived in a rural or specific minority area and sometimes on the size of the minority population. The fertility policies for ethnic minorities depended on the strategy the province uses toward those minorities. Autonomous regions generally had more lenient fertility constraints for minority couples. Most autonomous regions are not included in this study as the individual data set does not cover those

⁵The provinces are Hainan, Yunnan, Qinghai, Ningxia and Xinjiang. In the province of Guangdong, couples with a rural household status were also allowed to have two children until 1998 (Scharping (2013)).

regions.

Following a statement from the central government, all provinces at some point introduced the criterion that if one or both spouses were an only child they were eligible. This policy was motivated by the idea that the one-child-per-family policy should only hold for one generation. Provinces implemented this criterion over the course of the 1990s to early 2000.

There were specific exemptions for certain occupational groups such as fishermen and mine workers, as well as for veterans. Some provinces also had rules for couples who already adopted a child or that had their first child overseas. These policies are not taken into account here because they are not observable or can be anticipated at the age of 16. The category of couples with “real difficulties” is the vaguest and potentially flexible one. This makes an evaluation impossible without governmental application and acceptance data.

To summarize, I use the following criteria for the empirical analysis:⁶

1. In five provinces, couples in rural areas were allowed to have two children. In the province of Guangdong, second child permits were given to couples in rural areas until 1998.
2. Couples living in rural areas whose first child is a girl can anticipate being eligible with a likelihood of 50%.⁷
3. Couples in which one or both spouses belong to a national minority can have two children. This can apply to the whole province, the rural area or a specific minority area.
4. Couples in which one or both spouses are an only child can have two children.

1.3 Model

The motivation of the theoretical part is two-fold: First, it illustrates different channels through which anticipated fertility can affect educational investment in the most traceable form. Second, the model shows how different types of individuals react differently to a change in the cost of another child. Some anticipate to have another child, some do not change the number of children planned but benefit from not having to pay the fine. I use the results to guide the empirical investigation which aims at isolating group-specific effects. These two goals are achieved with a simple two-stage model, which can be extended for additional questions.

⁶Specific exemptions I also use: In Jiangsu province, men can have a second child if the first born is a girl and they do not have a brother. In Jilin province, in rural areas, if one spouse is an only child and the first born is a girl, the can have a second child.

⁷50% is an approximation. However, there are no reports of couples falling into that exemption that tried to avoid having a son as the first child. The main simplification is that I ignore sex differences in the cost-benefit analysis of parents.

1.3.1 Set-up

In the model, a representative family consists of two parents and their child, the individual who is a teenager at the beginning. They must decide how much to invest in the education of the teenager in the present (period 1). In the future (period 2), the individual is grown up, married, earns income together with their spouse. The newly formed couple can have children themselves. The parents in period 1 are potential grand-parents. The family has to pay a fine for the second child if they are not eligible for a second child permit. The fine depends on the educational level of the now grown-up individual.

The educational decision and fertility decision are modeled as made by the family. There are some argument for this is a simplifying assumption. For one, the teenager can influence educational investment by making more or less effort and by persuading the parents of their school choice. Parents can influence the teenager's fertility decision by passing on their own fertility preferences and by offering their help raising their grand-children. This is particularly persuasive in China where the family is still the most important social unit for many individuals. Also, grand-parents draw important benefits from having grand-children. They are invested in the continuation of their family line, besides other biological, social and altruistic motivations. However, they may see having grand-children as being particularly costly. While the the grown-up teenager raises the children, less resources are available for the then grandparents. This can overlap with the time when they are retired and have financial and care needs.

In the model, the educational investment stage and the reproductive stage do not overlap. This is realistic with regard to primary and secondary education in China, which are usually finished before starting the reproductive stage. In China, the minimum age for marriage is 20 for women and 22 for men and individuals are strongly discouraged from having children without being married.

Period 1 “the present”

In period 1, the family consumes c^1 . The income Y of the family is given exogenously and is used for consumption, investment in education I or saving s . Utility in period 1 is given by

$$u(c^1) = u(Y - \eta I - s) \tag{1.1}$$

where $u(\cdot)$ is the utility of consumption, assumed to be strictly increasing and concave. $\eta > 0$ is the constant cost of education, Y the exogenous income, I the educational level of the teenager and s the level of savings. I assume that families are credit constrained such that $s \geq 0$.

Period 2 “the future”

In period 2, the teenager is grown up and married. The income earned is consumed by the family

and spent on the n children the newly formed couple has. The family gets utility from having children which is I assume is additively separable from the utility of consumption.

Utility in period 2 is given by:

$$u(c^2) + \alpha_i h(n) = u(y(I, J, n) + Rs - p(I, n, Z) - f(n)) + \alpha h(n) \quad (1.2)$$

where $p(I, n, Z)$ indicates the fine that the family has to pay with $Z \in \{0, 1\}$ indicating eligibility status. The fine is dependent on the educational level of the teenager I . This reflects the fact that monetary fines are dependent on the household income and that the family might have to pay non-monetary fines such as losing the job or not being promoted. The fine is 0 for the first child and only has to be paid when the family is not eligible for a second child permit (i.e $Z = 0$):

$$p(I, n, Z) = \begin{cases} 0 & \text{if } n = 1 \\ 0 & \text{if } n = 2 \text{ \& } Z = 1 \\ \rho(I)(n - 1) & \text{if } n \geq 2 \text{ \& } Z = 0 \\ \rho(I)(n - 2) & \text{if } n \geq 2 \text{ \& } Z = 1 \end{cases} \quad (1.3)$$

$f(n)$ is the cost of raising n children (strictly increasing and concave), $\alpha h(n)$ represents the utility of having children (strictly increasing and concave), with α being an individual fertility preference parameter drawn from a given distribution. $y(I, J, n)$ is the household income, assumed to be strictly increasing and concave in the educational level I , the spouse's educational level J and decreasing in the number of children n . For now, we assume that the educational level of the spouse J is exogenous. J is endogenized in section 1.3.4. An example is an income generation function that remunerates individuals for each hour worked multiplied by their productivity : $y(I, n) = (T - \mu n)P(I)$. Productivity $P(I)$ is a concave function of education and T is the maximum time an individual can work. Working hours decrease μ for each child. One can interpret T as the number of years an individual works in their life where the number of children decreases the years of working.

Given their fertility preferences and eligibility status, family i solves:

$$\max_{I, s, n} u(Y - \eta I - s) + \delta[u(y(I, J, n) + Rs - p(I, n, Z) - f(n)) + \alpha h(n)] \quad (1.4)$$

with δ as the discount factor.

1.3.2 Optimal educational level and optimal number of children

Maximising with respect to education gives the optimal level of education as a function of the number of children n :

$$u'(c^2) \left[\frac{\partial y(I^*, J, n)}{\partial I^*} - \frac{\partial p(I, n, Z)}{\partial I^*} \right] = \frac{\eta}{\delta} u'(c^1) \quad (1.5)$$

The number of children affects educational investment by decreasing family consumption in the second period, by affecting the returns to education directly through $\frac{\partial y(I^*, J, n)}{\partial I^*}$ and indirectly through $\frac{\partial p(I, n, Z)}{\partial I^*}$ when the number of children is higher than 1.

Maximising utility with respect to n gives us the optimal number of children as a function of education:

$$u'(c^2) \left[\frac{\partial p(I, n^*, Z)}{\partial n^*} + \frac{\partial f(n^*)}{\partial n^*} - \frac{\partial y(I, J, n^*)}{\partial n^*} \right] = \alpha \frac{\partial h(n^*)}{\partial n^*} \quad (1.6)$$

On the left-hand side is the marginal (opportunity) cost of having n^* children. It consists of the marginal cost of raising and educating n^* children ($\frac{\partial f(n^*)}{\partial n^*}$), the fine if the family has more than one child and a decrease in income due to shorter working hours. On the right-hand side is the marginal benefits of having n^* children. The effect of education on the optimal number of children reflects standard results: Education increases income and thus makes having children less costly. However, education also increases the opportunity cost of having children through $\frac{\partial y(I, J, n^*)}{\partial n^*}$. Higher education also implies having to pay a higher fine for the second child, increasing the cost of having another one.

Maximising with respect to savings adds the following constraint:

$$R = \frac{u'(c^1)}{\delta u'(c^2)} \quad (1.7)$$

1.3.3 Effect of 2nd child permits

I focus my attention to the choice set for the number of children being either one or two ($n \in [1, 2]$). This encompasses the choice set of the majority of Chinese (the theoretical discussion easily extends to an unrestricted choice set). The number of children is a discrete variable and there is no fine for the first one. Therefore, eligibility does not necessarily change the optimal number of children given by equation 1.6. The effect of the exemption depends on if eligibility changes anticipated fertility or not. There are three cases (with the standard policy evaluation names in brackets):

1. **Always one child (never-takers):** $n^*(Z = 0) = 1$ and $n^*(Z = 1) = 1$

The optimal number of children with eligibility or without is the same. The family is unaf-

ected by the exemption.

2. **Always two children (always-takers):** $n^*(Z = 0) = 2$ and $n^*(Z = 1) = 2$

The fertility decision is not altered by the eligibility status because the family always wanted two children. But the family benefits from not having to pay the fine for the second child.

3. **Increasesers (compliers):** $n^*(Z = 0) = 1$ and $n^*(Z = 1) = 2$

Becoming eligible, the optimal number of children increases by one child.

The model disregards any level of uncertainty and assumes full information. To illustrate the basic mechanisms, this model suffices. As the model does not have uncertainty, it is known in which category the family will fall.⁸

Second child permits affect the educational investment of *Increasesers* and *Always-2-children* as followed:

Always 2 children:

Optimal education with $Z = 0$:

$$u'(y(I^*, J, 2) + Rs - \rho(I^*) - f(2)) \left[\frac{\partial y(I^*, J, 2)}{\partial I^*} - \frac{\partial \rho(I^*)}{\partial I^*} \right] = \frac{\eta}{\delta} u'(Y - \eta I^* - s) \quad (1.8)$$

Optimal education with $Z = 1$:

$$u'(Y(I^*, J, 2) + Rs - f(2)) \left[\frac{\partial y(I^*, J, 2)}{\partial I^*} \right] = \frac{\eta}{\delta} u'(Y - \eta I^* - s) \quad (1.9)$$

Being eligible has two opposing effects:

Intertemporal consumption smoothing: Eligibility decreases the marginal utility of consumption in period 2 ($u'(c^2)$) because the fine $\rho(I)$ does not have to be paid any more. This decreases education because the family responds with higher consumption and lower educational investment in period 1.

Returns to education: Eligibility increases the returns to education by $\frac{\partial \rho(I^*)}{\partial I^*}$ because the fine is dependent on income which increases in education.

Since there are two potential opposing effects, the total effect can be negative or positive.

Increasesers:

⁸One could easily include for example uncertainty in the income gained in period 2 by adding a random shock to the income term. Given the distribution of this term, one knows the distribution in the type of family. For example, a family with a given α could an *increaser* with 80% likelihood and *always-2* with 20% likelihood.

Optimal education with $Z = 0$:

$$u'(y(I^*, J, 1) + Rs - f(1)) \left[\frac{\partial y(I^*, J, 1)}{\partial I^*} \right] = \frac{\eta}{\delta} u'(Y - \eta I^* - s) \quad (1.10)$$

Optimal education with $Z = 1$:

$$u'(y(I^*, J, 2) + Rs - f(2)) \left[\frac{\partial y(I^*, J, 2)}{\partial I^*} \right] = \frac{s}{\delta} u'(Y - \eta I^* - s) \quad (1.11)$$

Again, eligibility has two effects:

Intertemporal consumption smoothing: Eligibility increases the marginal utility of consumption in period 2 ($u'(c^2)$) by increasing spending on children by $f(2) - f(1)$. Also, the individual earns less due to having to care for two children (when $y(I^*, J, 1) > y(I^*, J, 2)$). Marginal utility of additional earning in the future increases and the family uses education as a way to shift consumption from period 1 to period 2, such that equation 1.5 holds.

Labour supply and returns to education: Eligibility decreases the returns to education if the individual has to cut productive working hours ($\frac{\partial y(I^*, J, 2)}{\partial I^*} < \frac{\partial y(I^*, J, 1)}{\partial I^*}$). This decreases returns to education and thus decreases the incentives to invest in education.

The sign of the overall effect is discussed in the appendix in section 1.9.3.1 using as an example the income generation function $y(I, n) = (T - \mu n)P(I)$. The effect is positive if the intertemporal consumption smoothing effect is stronger than the labour supply effect.

When the cost of children increases, this strengthens the income smoothing channel. However, it also enforces the labour supply effect as the lost income due to a reduction in working hours is more painful. If there is no labour supply channel, higher cost of raising children unambiguously increases the effect of anticipate fertility on education. Also, if we assume constant absolute risk aversion (CARA), higher costs of raising children always increases the effect.

I also discuss the effect of the loss in working time due to child care in the appendix (section 1.9.3.3). A higher loss in working time increases the labour supply channel and has an ambiguous effect on the income channel. If we assume CARA, less working time lost due to a child (i.e. a the smaller μ) increases the effect of anticipated fertility on education.

These factors are summarized in the following proposition:

Proposition 1 *Assuming, CARA, the overall effect of anticipating another child on education:*

1. *increases in the cost of raising a child $f(n)$,*
2. *decreases in the loss of working time due to childcare δ .*

Proofs are in the appendix in section 1.9.3.3 and 1.9.3.2.⁹

1.3.3.1 The type of family according to fertility preferences α

The effect of the policy depends on if the family is behaving as an *always-2*, *always-1*, or *increaser*. It depends on the exogenous parameters α and gives the following comparative statics: For given cost and fine levels, families with a low fertility preference α are *always-1*, those with a medium α are *increasers* and those with a high α are *always-2*. This is illustrated in figure 1.4 in the appendix. Define the threshold $\underline{\alpha}$ as the α at which for $Z = 1$, equation 1.6 is equal for $n = 1$ and $n = 2$. Define the threshold $\bar{\alpha}$ as the α at which at $Z = 0$, equation 1.6 is equal for $n = 1$ and $n = 2$.

Proposition 2 *If*

1. $\alpha < \underline{\alpha}$: the family has 1 child independent of eligibility (*always-1*).
2. $\alpha > \underline{\alpha}$ and $\alpha < \bar{\alpha}$: the family has one child if not eligible and two children if eligible (*increaser*).
3. $\alpha > \bar{\alpha}$: the family has two children in any case (*always-2*).

This follows from the fact that the number of children n^* is increasing in α and by the definition of the thresholds. More interestingly, how do these threshold change according to other variables?

Proposition 3 *The threshold $\bar{\alpha}$, above which the family is an always-2,*

- *increases in the cost of raising a child $f(\cdot)$,*
- *increases in fines $\rho(\cdot)$.*

The threshold $\underline{\alpha}$, above which the family is an increaser

- *increases in the cost of raising a child $f(\cdot)$,*
- *is unaffected by the fines $\rho(\cdot)$.*

In section 1.9.3.5 in the appendix, the the equations for $\underline{\alpha}$ and $\bar{\alpha}$ are derived. $\underline{\alpha}$ does not depend on the fines, while $\bar{\alpha}$ does. The cost of raising children $f(\cdot)$ increases the total cost of children and decreases the number of children for a given α .¹⁰ Being indifferent between having 1 or 2 children,

⁹Furthermore, assume the hourly wage depends on productivity $P(I)$ multiplied by a factor τ which captures technology or labour market conditions. Assuming CARA, the the effect of anticipated fertility on education increases with τ . This is shown in section 1.9.3.4.

¹⁰Let's reformulate $f(n)$ as a concave function $f(n)$ scaled by the parameter ψ . Equation 1.6 is rewritten as $u'(c^2) \left[\frac{\partial p(I, n^*, Z)}{\partial n^*} + \frac{\psi \partial f(n^*)}{\partial n^*} - \frac{\partial y(I, J, n^*)}{\partial n^*} \right] = \alpha \frac{\partial h(n^*)}{\partial n^*}$. Taking total derivatives of the left hand side for ψ gives: $u'(c^2)\psi + u''(c^2)(-1) \frac{\partial f(n^*)}{\partial n^*} \frac{\psi \partial f(n^*)}{\partial n^*} > 0$.

conditioned on eligibility ($Z = 1$) implies higher fertility preferences when the cost of raising children is higher. The same hold for being indifferent between having 1 or 2 children conditioned on not being eligible ($Z = 0$). As fines are essentially an increase in the cost of having another child, $\bar{\alpha}$ increases in the level of the fines.

Second child permits should not have any effect on the always-1. A positive or negative overall effect can be driven by families that are *benefiters*, or *always-2*, or both of them. Proposition 2 helps to empirically disentangle the overall effect of second child permits on schooling investment, by differentiating between the effect it has on individuals that are *increasers* and *always-2*.

1.3.3.2 Effect on saving

In this model, saving and educational investment are both tools to smooth consumption. Denote E_I the net returns to education $\frac{1}{\eta} \left[\frac{\partial y(I^*, J, n)}{\partial I^*} - \frac{\partial p(I, n, Z)}{\partial I^*} \right]$, such that the first order condition 1.5 can be rewritten as:

$$E_{I^*} = \frac{u'(c^1)}{\delta u'(c^2)} \quad (1.12)$$

The family uses only education if at $I = I^*$ given by equation 1.5 $E_{I^*} > R$ (case i). It uses only savings if at $s = s^*$ given by equation 1.7 $E_{I^*} < R$ (case ii). It uses both if at $I = I^*$ and $s = s^*$ $E_{I^*} = R$ (case iii).

Intertemporal consumption smoothing: If with anticipated fertility the need for intertemporal consumption smoothing increases, in case i, the family uses educational investment for smoothing until either $E_{I^*} = \frac{u'(c^1)}{\delta u'(c^2)}$ or $E_{I^*} = R$. In case ii and iii, because the returns to savings are linear and the returns to education are decreasing, the family only uses savings to smooth. However, under certain assumptions, the family increases investment in education and savings as a reaction to increased anticipated fertility. For example, when educational investment is considered safe and savings are considered risky (or the other way around).

Labour supply effect: Higher anticipated fertility decreases the expected labour supply and thus the returns to education. This can tilt the favour towards using savings rather than educational investment for consumption smoothing.

1.3.4 Gender differences and marriage

In this section, I distinguished between families with male or female teenager. I also consider marriage market returns to education. First, I assume that spouses match positive-assortatively on educational levels. The spouse's educational level J is a function of own education I : $J(I) = \sigma I$ with $0 < \sigma \leq 1$. This captures the correlation between the educational levels. Furthermore, I assume that only women reduce their working hours when they have a child and that the spouses'

incomes enter additively. Both assumptions are in line with descriptive statistics (see section 3.5.4).

The maximisation problem of a family with a daughter is:

$$\max_{I,s,n} u(Y - \eta I - s) + \delta[u(y(I, n) + y(J(I)) + Rs - p(I, n, Z) - f(n)) + \alpha h(n)] \quad (1.13)$$

$$u'(c^2) \left[\frac{\partial y(I^*, n)}{\partial I^*} + \frac{\partial y(J(I^*))}{\partial J(I^*)} \sigma - \frac{\partial p(I, n, Z)}{\partial I^*} \right] = \frac{\eta}{\delta} u'(c^1) \quad (1.14)$$

Families with a daughter who anticipates a reduction in working hours expect to experience the *labour supply effect*. However, this effect gets mitigated by the returns to education in the marriage market $\frac{\partial y(J(I^*))}{\partial J(I^*)} \sigma$. While the labour market returns to education $\frac{\partial y(I^*, n)}{\partial I^*}$ depend on the number of children, the marriage market returns do not. The stronger the correlation of educational levels, the better the daughter's educational investment can be used for intertemporal consumption smoothing. For families with sons, it is the opposite. Though the marriage market returns to education are also positive, they decrease in the anticipated number of children.

However, for men, there might be another factor, relevant in this context. So far, I assume that everyone marries. I now introduce the possibility of staying single. If a person stays single in period 2, he/she earns income and does not have children: $u(c_s^2) = \tilde{y}(I)$. Denote c_m^2 the consumption in period 1 when being married and c_s^2 the consumption when staying single. Denote ω the probability of marrying in period 2 and τ the ratio of men to women. I assume that the marriage surplus, defined as the utility of being married minus the utility of staying single, is positive independent of the educational level of the individual and the potential spouse and the number of children: $u(y(I) + y(J(I), n) + Rs - p(I, n, Z) - f(n)) + \alpha h(n) - u(\tilde{y}(I)) > 0, \forall I, n$. As a consequence, if there are as many men as women, then everyone gets married in period 2. If the sex ratio is skewed, all of those of the scarce sex get married and some of those of the abundant sex stay single. I assume that the likelihood of marriage for the abundant sex is a function of the sex ratio and the individual's education: $\omega(I, \tau)$.¹¹ We focus on the situation where there are more men than women, which is the relevant case for China: $\tau > 1$. The utility of a man is then given by (FOC in section 1.9.3.6 in the appendix):

$$\max_{I,s,n} u(Y - \eta I - s) + \delta \left[\underbrace{\omega(I, \tau) u(y(I) + y(J(I), n) + Rs - p(I, n, Z) - f(n)) + \alpha h(n)}_{\text{married}} + (1 - \omega(I, \tau)) \underbrace{u(\tilde{y}(I))}_{\text{single}} \right] \quad (1.15)$$

¹¹The assumption is related to assuming non-transferable utility. In the transferable utility case, families could compensate a lack of education with monetary transfers. By restricting monetary transfers, through transaction cost, limited commitment or simply social norms, families have to use pre-marital investment such as education in order to make their child attractive in the marriage market.

Marriage market effect: An increase in the number of anticipated children implies a (weak) increase in the marriage surplus, because it increases the utility of being married but not the utility of being single. For *always-2-children*, it increases the utility of marriage as it takes away the monetary fines. For *increasers*, it (weakly) increases the utility of being married as they choose to have two children over one child because it given them higher utility. It increases the incentives to invest in education to increase the odds of finding a spouse to make the now more beneficial situation more likely. The marriage market effect are thus be positive for *increasers* and *always-2*, but not necessarily of the same magnitude.

Proposition 4 *An increase in anticipated fertility increases the incentives to invest in education to improve the odds of finding a spouse, given that the sex ratio is skewed.*

As in the previous sections, the general equilibrium effects are not taken into account. The likelihood of finding a spouse is not only dependent on own education but also the distribution of educational levels within the same marriage market. If several individuals within the same marriage market become eligible, this can change the educational distribution of the abundant sex.

1.3.5 The effect of fines on education

How does the amount of fines change the incentives to educate? Second child permits can be interpreted as a reduction in the cost of having a second child, and the size of the cost reduction depends on $p(I)$. The effect of these monetary fines is again different for different types of individuals. For *increasers*, defined as those that anticipate a second child only if they are eligible for a second child permit, the size of the fine does not matter. However, as mentioned in proposition 2, how many *increasers* there are in the population depends on the size of fine. The higher the monetary fines, the higher the cost reduction due to 2nd child permits, thus the more people react to it.

On top of changing the composition in the population between *increaser* and *always-2*, the amount of the fines change the incentives to educate for *always-2*. Higher fines, as a function of household income, increase the need to smooth consumption intertemporally, but also have a stronger negative effect on education. Higher fines also strengthen the marriage market effect: not having to pay high fines increases the marriage surplus more than not having to pay low fines.

1.4 Data

1.4.1 Individual Data

For the empirical analysis, I use individual survey data from the 2010 China Family Panel Study (CFPS). It was designed by a Peking University research team, supported by Peking University

985 funds and carried out by the Institute of Social Science Survey of the Peking University. The dataset is available online in English and Chinese. Not all provinces are represented in the sample. In particular, the sample does not cover the autonomous regions of China except for Guangxi Zhuang. The main population is sampled from Gansu (12%), Henan (11%), Guangdong (9%), Shanghai (8.5%) and Liaoning (8.4%).

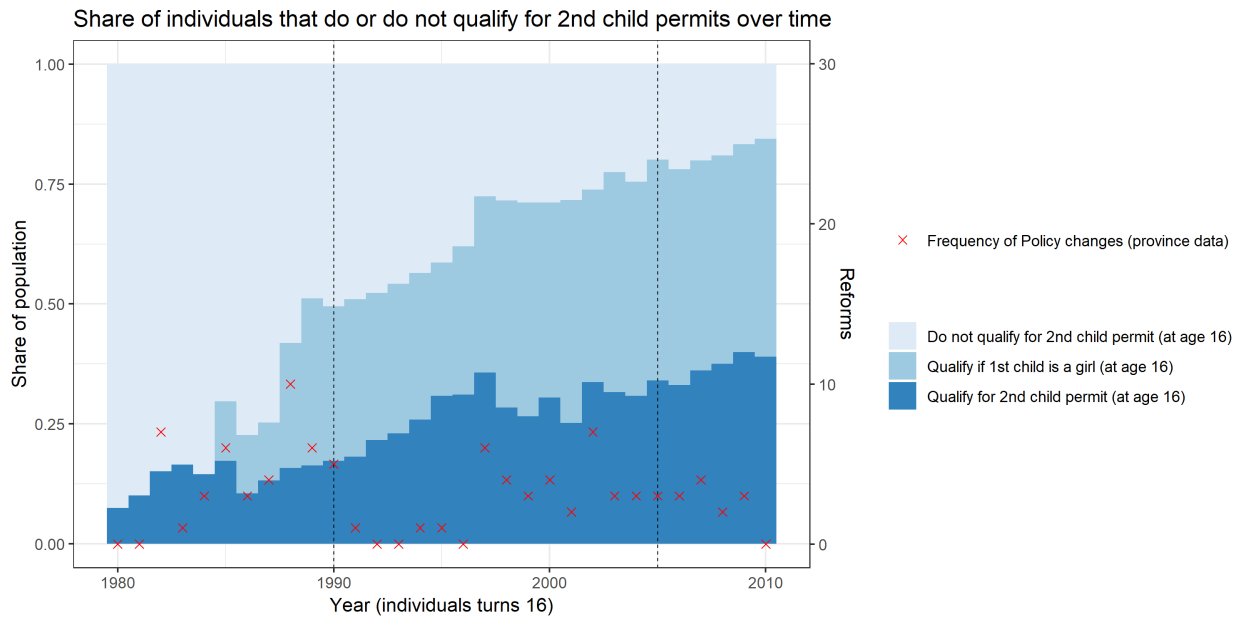
For the main sample, I use individuals that turned 16 between 1990 and 2005 which leads to a sample of 7840 observations of which 53% are female. Summary statistics are displayed in table 1.6 in the appendix. The sample is predominantly rural: 70% hold an agricultural household status and 30% hold a non-agricultural household status. 89% of the sample indicate that they are of majority Han ethnicity. The main ethnic minorities in the sample are Miao (2.1%), Yi(2.3%) and Man (1.5%). On average, individuals stayed in school for 7.6 years (women: 7.1, men: 8.2). Those with a non-agricultural household status spend nearly twice as many years at school as those with agricultural household status (11.4 compared to 6.1).

Empirical identification relies on geographical and temporal variation in the eligibility criteria. Figure 1.2 illustrates the share of individuals that fulfilled an eligibility criterion at the age of 16. I differentiate between two categories: those that fulfil a criterion to have another child, and those that fulfil the criterion to have a second child if the first child is a girl (i.e. being eligible with a likelihood of approximately 50%). There was high variation at the beginning of the OCP between 1982 and 1990 when provinces introduced criteria for the first time. However, the variation is too high to allow individuals to form adequate expectations. In the time span between 1990 and 2005, there are some policy reforms are provinces update their criteria catalogues, but the policy setting is mostly stable.

Those in the main sample are too young to have finished their reproductive stage at the time of the survey. I am thus unable to use the number of children allowed at secondary school age as an instrument for the actual number of children. I, therefore, use individuals that turned 16 between 1980 and 1995 as the older sample because they have already finished their main reproductive stage. With the older sample, I can investigate the effect of second child permits on the likelihood of having a second child and the treatment heterogeneity. Summary statistics for the older sample are also displayed in table 1.6. As expected, educational levels are lower but other characteristics are the same (sex, ethnicity). There are more individuals with an agricultural household registration status.

1.4.2 Policy exemptions

The information about province level policies is compiled from [Scharping \(2013\)](#), [Baochang et al. \(2007\)](#) and official family planning regulation documents. The documents were accessed



Data Source for individual data: CFPS 2010. Figure shows if an individual qualifies for a second child permit the year they turn 16. Includes individuals that turn 16 between 1980 and 2010. Frequency of reform changes is based on province level data. Source: [Scharping \(2013\)](#), [Baochang et al. \(2007\)](#) and family planning regulations. One reform is defined as a change for one sub-group.

Figure 1.2: Share of population qualifying for 2nd child permit and incidences of second child permit reforms.

online in Mandarin Chinese and translated into English.¹² An excerpt of the data is displayed in table 1.5 in the appendix. Based on this information, I calculate if at the time of educational decision-making an individual fulfils a second child permit criterion. The main educational decision is made around the age of 16. At this time, students on average must decide whether to continue with senior general or vocational high school after completing compulsory junior high school.

For the main analysis, I use those who turned 16 between 1990 and 2005. This has several reasons: First, only after 1990 Chinese citizens had official legal documents that they could rely on. Before, conditions for second child permits were only presented as guideline. It is debatable whether implementation and knowledge of exemptions were comparable between provinces. Second, I do not want to mix up the effects of the policy with the implementation time of the 1986 compulsory secondary school reform. To assure that in the year of the survey (2010), basic educational investment is finished, the upper threshold is set at 2005. Also, as previously mentioned, in this time span there are several reforms in different provinces but not at the same time. Yet, the policy framework is sufficiently stable so that fertility anticipations can be formed based on the current policy situation.

1.5 Empirical methodology

The goal of this study is to isolate the effect of anticipated fertility on education. I denote I the educational outcome and \tilde{N} anticipated fertility. Furthermore, denote with P the cost of an out-of-quota child. Ideally, one would want to estimate:

$$I_i = \theta_1 \tilde{N}_i + \theta_2 P_i + \mu_i \quad (1.16)$$

using anticipated eligibility $\tilde{Z} \in \{0, 0.5, 1\}$ as an instrument for \tilde{N} . The first stage would take the following form:

$$\tilde{N}_i = \beta \tilde{Z}_i + X_i + \epsilon_i \quad (1.17)$$

where X indicated individual-level characteristics that determine eligibility status, province fixed effects and time fixed effects. Conditioned on X , the anticipated eligibility status \tilde{Z} is assumed to be uncorrelated with the error term ϵ .

1.5.1 Effect of eligibility on education

However, \tilde{N} and P are unobserved. Instead, only the actual fertility outcome N is observed in the older sample for which the actual eligibility status $Z \in \{0, 1\}$ is known after the birth of the first child. So what can be identified? Denote $I(\tilde{Z})|X$ the educational level dependent on anticipated eligibility at the age of 16 given characteristics X . I can identify the intention-to-treat

¹²This collection of policy information collected by Wanying Zhao is available upon request in English.

effect $E(I(1)|X) - E(I(0)|X)$ which compares the educational level of individuals that fulfil a second child permit criterion at secondary school age with those that are not, with the following regression:

$$I_i = \tilde{\theta}\tilde{Z}_i + X_i + \tilde{\mu} \quad (1.18)$$

The intention-to-treat effect $\tilde{\theta}$, estimated in section 1.6.1, is a policy-relevant parameter. From theory we know that it is a weighted average of the effect of the policy on the three different groups *always-1*, *increasers* and *always-2*, which correspond to never-takers, compliers and always-takers (Angrist et al. (1996)). Other fertility policies, such as child benefits or paid parental leave, that change the costs of a child, would also affect compliers and always-takers. The intention-to-treat effect is thus a parameter of interest, as it indicates how the policy overall affects education. However, it is also of interest to understand how the two different groups react, in order to be able to make policy recommendations for other settings where the composition of the population is different.

In order to interpret the coefficient as the effect of the policy on compliers, one has to assume that the behaviour of the other groups is unaffected (exclusion restriction). In this context, it is easy to do for the group of *always-1*, but it is restrictive in the case of *always-2* since, as we have seen in the theory section, they benefit from not having to pay income-dependent fines.

The intention-to-treat effect can be decomposed into the effect on *increasers* which goes through anticipated fertility \tilde{N} and on *always-2* whose effect goes through the fine for an out-of-quota child. By definition, *increasers* are those that change anticipated fertility from 1 to 2 as a response to eligibility, thus β is the share of *increasers* in the population. Denote π the share of *always-2* in the population and ρ the average amount of fines. The average effect of eligibility on out-of-quota fines then defined as: $E(P(1)|X) - E(P(0)|X) = \pi\rho$. Thus, the intention-to-treat effect can be decomposed into:

$$E(I1|X) - E(I(0)|X) = \theta_1[E(\tilde{N}(1)|X) - E(\tilde{N}(0)|X)] + \theta_2[E(P(1)|X) - E(P(0)|X)] \quad (1.19)$$

$$= \beta\theta_1 + \pi\rho\theta_2 \quad (1.20)$$

In the standard LATE case, the exclusion restriction implies that $\pi\rho\theta_2 = 0$: either there are no always-takers or always-takers are unaffected by the instrument Z . Furthermore, if anticipated fertility \tilde{N} was observed, one can estimate β , the “take-up rate”. Even if the take-up rate is unobserved, by assuming a non-negative effect of the instrument on take-up together with the exclusion restriction, one is able to identify the sign of the effect. Yet, without the exclusion restriction, this is not possible.

1.5.2 Effect of eligibility on fertility outcome

Since anticipated fertility is unobserved, β cannot be estimated. However, it is reasonable to assume that the anticipated eligibility status can only affect anticipated fertility when eligibility

affects the fertility outcome. Fertility outcome N is observed in the older sample, as is eligibility status $Z \in \{0, 1\}$ after the birth of the first child, so that we can estimate:

$$N = \tilde{\beta}Z + X + \epsilon \tag{1.21}$$

This regression provides an estimator of the take-up rate of the 2nd child permit categories included in Z . Results are discussed in section 1.6.2. One should keep in mind that there are criteria for 2nd child permits that are not included in Z because they are not observable or because they rely on the realization of the first birth or occupational choice. Therefore, $\tilde{\beta}$ is not the take-up rate of all 2nd child permits criteria. However, $\tilde{\beta}$ measures the effect of categories important for forming fertility anticipations.

1.5.3 Isolating the effect on *increasers*

In the older sample, we observe individuals that have a second child without being eligible. Keeping in mind the caveat that some of those were eligible for different 2nd child permits, we can nonetheless use the information gained from this sample to decompose the overall intention-to-treat effect, under some assumptions, guided by the theoretical model.

In the theoretical model, fertility preferences determine if the individual is an *always-1*, *increaser*, or *always-2*. The share of *always-2* is denoted with π . I use the fertility proxies W to predict if an individual in the main sample will be a *always-2* or not. For this, those not eligible in the older sample are used to estimate the effect of W on the likelihood of having a second child:

$$Prob(N_i = 2|Z_i = 0, X) = f(\alpha W_i|Z_i = 0, X) \tag{1.22}$$

The purpose of this exercise is prediction and not inference, thus, the values of α are not of interest and only capture correlation, not necessarily causality. To get the highest predictive power and to account for potential non-linearities, I use a classification random forest with 1000 trees, where the prediction is determined by majority rule at the respective node. The number of predictors used for each tree is determined by minimising the out-of-bag error rate.

The effect of the fertility predictors, based on the training set of those in the older sample who are not eligible, is then extrapolated to the whole older sample and the main sample. Denote the prediction status of an observation L which can be either 0 (not predicted to be an *always-2*) or 1 (predicted to be an *always-2*). Regressions 1.26 and 1.18 are then run with L as an interaction term. We therefore observe the coefficients $\tilde{\theta}(L) = E(Y(1) - Y(0)|L)$.

Denote the share of *always-2* in the sample of those predicted to be *always-2* $\bar{\pi}$ and the share in the sample of those predicted to not be *always-2* $\underline{\pi}$. If the prediction is informative about the underlying distribution of *always-2*, then $\bar{\pi} > \pi > \underline{\pi}$. Under the assumption that the prediction

is informative, the decomposition allows to make statements about the signs of the coefficients θ_1 and θ_2 .

It is again useful to draw comparisons to the standard case in which \tilde{N} is observed. In this case, we could estimate the share of *always-2* (π) and *increasers* (β). Once they are known this decomposition allows to estimate θ_1 and θ_2 separately (as shown in Hull (2018)).

1.6 Results

1.6.1 Effect of second child permits on education

The first and most relevant educational choice for most families is made when the teenage child is 16 years old. This is the age at which I calculate anticipated eligibility. Educational attainment is first measured in years of education ($I_{ip\ in\ 2010}$). It is then replaced by a dummy variable indicating enrolment into senior high school and college or university. I estimate 3 different specifications. Each specification relaxes an assumption but also includes more control variables and thus puts more burden on the model. Specification 1 is a DID estimation that assumes a national trend (equation 1.23). It includes year fixed effects (now denoted γ), province-specific effects (denoted κ), characteristics defining potential eligibility groups and additional controls (X). It assumes that annual effects are the same for each province and each population.

$$I_{ip\ in\ 2010} = \tilde{\theta}Z_{ip\ age(16)} + X'_i\kappa_p + \gamma_t + \kappa_p + \epsilon_{ipt} \quad (1.23)$$

Specification 2 is a DID estimation with province-specific trends (equation 1.24). It relaxes the assumption that the annual effects are the same for each province. It still assumes that the annual effects are the same for each population within the province.

$$I_{ip\ in\ 2010} = \tilde{\theta}Z_{ip\ age(16)} + X'_i\kappa_p + \gamma'_t\kappa_p + \epsilon_{ipt} \quad (1.24)$$

The third specification is a Differences-in-Differences-in-Differences (Double DID or Triple Differences approach, equation 1.25). It relaxes the assumption that the potentially eligible populations and the general population have the same trend.

$$I_{ip\ in\ 2010} = \tilde{\theta}Z_{ip\ age(16)} + X'_i\kappa_p + X'_i\kappa'_p\gamma_t + \gamma'_t\kappa_p + \epsilon_{ipt} \quad (1.25)$$

In specification 1 and 2, I compare the difference between those that were and those that were not eligible at age 16 within the same population in a province with a population where no change in eligibility occurs. In specification 3, I compare this difference-in-difference with the same comparison in another province where in the same time span no reform change for these populations

happened.

Identification comes from the geographical and time variation in the introduction and scope of second child permits. The double DID approach allows controlling for population-specific time trends. It thus relaxes the standard common trend assumption. The identifying assumption is that, conditioned on province trends, the sub-groups that become eligible have the same educational trend as the sub-groups in other provinces that do not change eligibility status (conditional group-specific common trend assumption). To support this identifying assumption, I run a pre-OCP placebo test. Identification also implies that provincial family planning policies are independent of educational measures that target the same group. For instance, if provinces that allow second child permits for ethnic minority couples increase these measures with an increase in the educational budget for ethnic minorities areas, the policy measure captures both. So far, I have not encountered evidence in the literature for such behaviour. I run additional robustness checks to verify that the results are not driven by a specific easily targeted group.

One might be concerned about potential spill-over through migration. However, the Chinese household registration system restricts the possibility to migrate particularly between provinces and between urban and rural areas. Applications for the second child permit can only be submitted at the place of registration. Moving the place of registration is difficult.¹³ Recent reforms are supposed to loosen these restrictions.

The results for the effect of second child permits on schooling investment are shown in table 1.1 with the specification 1, 2 and 3 in column 1, 2 and 3. Fulfilling a second child permit criterion at the age of 16 increases the years of education by around 0.9 years on average. The coefficient does not change significantly between the three specifications. Yet, standard errors increase. For the double Difference-in-Differences approach the coefficient is significant at 10% level with a p-value of 0.06. The coefficient is robust to controlling for the local fertility rate and the local sex ratio.

In table 1.2, I include the interaction with a female indicator and look at the effect on men and women separately. The effect is higher for men but not significantly so (see column 1). The coefficient for women alone lies at 0.6 years of education and is not significant. This empirical result suggests that men and to a smaller extent women who expect to be allowed to have two children increase their educational investment.

At which stage of the educational career does eligibility play a role? Instead of years of education, I now use indicator variables for finishing junior high school and enrolment into senior high school, technical college and university as outcomes. Results are displayed in table 1.7 in the

¹³Within the main sample, only 1.2% indicated a different provincial code as place of residence at the age of 12 than at the age of 3 while 5.4% indicated a different county or district code (within-province migration).

Dependent variable: Years of education					
	DID national trend (1)	DID province trends (2)	Double DID (3)	Double DID (4)	Double DID (5)
Anticipated Eligibility at Age 16	0.829** (0.385)	0.911** (0.440)	0.895* (0.467)	0.834* (0.442)	0.895* (0.467)
Local Fertility rate				-1.961*** (0.217)	
Local sex ratio (std)					0.0314 (0.186)
Female Dummy	Yes	Yes	Yes	Yes	Yes
Eligibility Indicators	Yes	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes	Yes
Eligibility Controls x Province FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Eligibility Controls x Year FE	No	No	Yes	Yes	Yes
Year FE x Province FE	No	Yes	Yes	Yes	Yes
Observations	7840	7840	7840	7840	7840
R^2	0.432	0.461	0.467	0.467	0.474

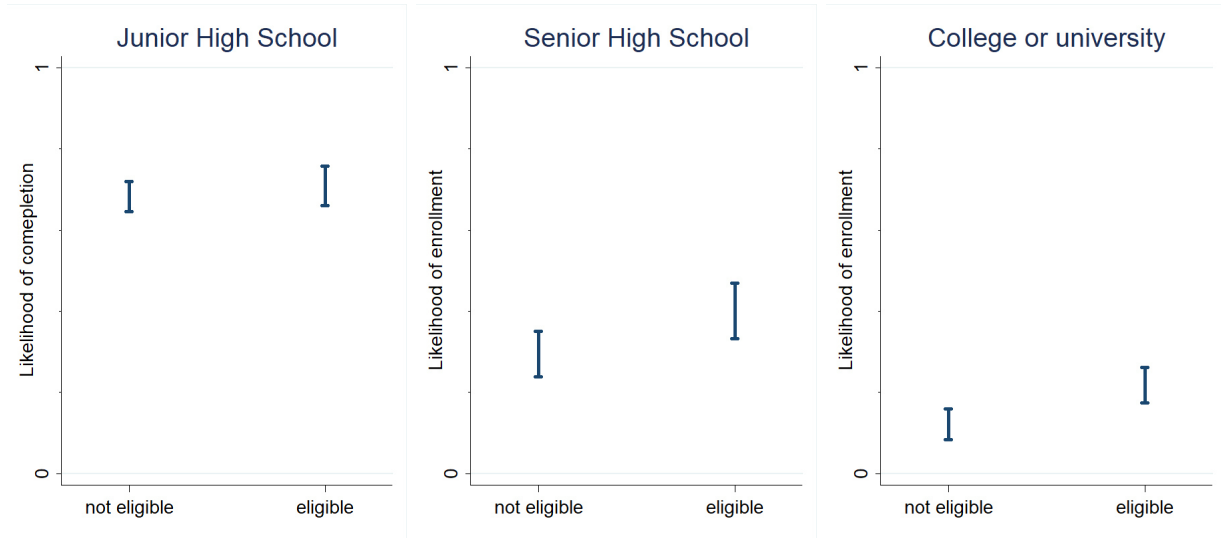
Note: Sample includes individuals that turned 16 between 1990 and 2005. The dependent variable is the years of education the individual completed in 2010. Standard OLS regression with standard errors clustered on provinces-times-urban-area levels in parenthesis. Significance levels: * 0.10; ** 0.05; *** 0.01. Anticipated eligibility at age 16 is calculated on the basis of the province, eligibility characteristics and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. Data source: China Family Panel Survey 2010.

Table 1.1: Effect of eligibility at age 16 on the years of education.

Dependent variable: Years of education			
	Double DID All (1)	Double DID Women Only (2)	Double DID Men only (3)
Anticip. Eligibility at 16	0.978* (0.494)	0.603 (0.582)	0.866* (0.474)
Female	-0.495 (0.539)		
Female X Anticip. Eligibility at 16	-0.194 (0.250)		
Double DID controls	Yes	Yes	Yes
Female X Double DID controls	Yes	No	No
Observations	7840	4184	3656
R^2	0.477	0.523	0.475

Note: Double difference controls: year fixed effects, province fixed effect, province times year fixed effects, eligibility controls, eligibility controls times year fixed effects and eligibility times province fixed effects. Sample includes individuals that turned 16 between 1990 and 2005. Column (2) includes only women, column (3) only men. The dependent variable is the years of education the individual completed in 2010. Standard OLS regression with standard errors clustered on provinces-times-urban-area levels in parenthesis. Significance levels: * 0.10; ** 0.05; *** 0.01. Anticipated eligibility at age 16 is calculated on the basis of the province, eligibility characteristics and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. Local sex ration: Sex ratio in the same district for those being born the same year plus and minus 3 years. Local fertility rate: Number of children individuals have in the same district that are between 30 and 35 years old in the year the individual turns 16. Data source: China Family Panel Survey 2010.

Table 1.2: Sex difference in the effect of eligibility at age 16 on the years of education.



Note: Displays 95% confidence intervals of marginal effects of eligibility. Cased on DID estimates with province specific trends (table 1.7). Estimates for not eligible and eligible (0.5 - eligible category is omitted)

Figure 1.3: Effect of anticipated eligibility on finishing junior high school, enrollment into senior high school and enrollment into college or university

appendix and graphically shown in figure 1.3. First, there is no effect of eligibility on junior high school completion (column 1 and 2). This is expected and can be regarded as a falsification test. At the age of 16, individuals have completed or are about to complete their mandatory junior high school degree. Second, enrolment into senior high school increases by around 9 percentage points (column 3 and 4). The effect carries on. Enrolment into at least college (vocational or university studies - column 5 and 6) increases significantly by around the same size. Enrolment into university increases by 6 percentage points (column 7 and 8).

Following, I run a placebo test to argue that the exclusion restriction is not violated (table 1.8 in the appendix). I map the eligibility status of those in my main sample to those that turned 16 before the introduction of any second child exemptions. I find that the placebo variable does not have any effect on the educational level of those that turned 16 between 1972 and 1982.¹⁴

As another robustness test, I check the effect of eligibility on the likelihood of senior high school enrollment at different age thresholds. If eligibility was simply correlated with provincial educational investment or local educational demand factors, there is no reason why the age cut-off of 16 should be relevant. We would expect similar results at the age of 17 or 18. However, if it is the second child permit reforms that drive the results, we should not see any effect on those who become eligible at the age of 17 or 18, because the decision to continue schooling or not has already been taken. Those that become eligible before the age of 16, should also be affected, but the effect

¹⁴Results do not change if I choose one year later or earlier.

should be smaller because those who become eligible at 16 would fall into the control group. The results are graphically displayed in figure 1.5 in the appendix. They are in line with the predictions if the effect was indeed driven by second child permit reforms. At age 17 and 18, the effects are insignificant and close to 0. At age 14 and 15, the effect size is only half as much as at the age of 16 and insignificant for the age of 14.

Furthermore, I investigate which group drives the results (see table 1.9 in the appendix). First, the coefficient for eligibility with 50% likelihood is not significantly different from zero. However, it is positive and within the expected range (column 1). I then include interaction terms for the different eligibility groups. The effect is the same for the minority and agricultural sub-groups (column 2 and 4). It is only significantly different for those who are an only child (column 3). It seems that anticipated eligibility does not have a significant effect on an only child. This can have several reasons. For one, those who are an only child might have low fertility preferences and thus fall into the category of *always-1*. Furthermore, the policy that allows only children to have two children was announced by the central government and slowly implemented in all provinces. It is possible that only children who were not eligible when they were 16 anticipated to be eligible once they were older.

Some criteria have to be fulfilled by both spouses. I check if this makes a difference to the effect of eligibility in table 1.10 in the appendix. Most criteria for rural households are conditioned on both spouses having a rural household status, nearly half of all eligible minorities fall into a spouse-dependent group and all of the only children. There is no significant difference between those that fulfil a spouse-dependent criterion and those that fulfil a spouse-independent criterion. The fact that all the only children have to find a spouse who is also an only child to be able to have two children offers another explanation why there is no effect on only children. They might not expect to find a partner who is an only child as well, as at this time only children still make up a small share of the population. On the contrary, it might not pose a burden to those of ethnic minorities or with a rural household status as the likelihood that they marry someone with the same status is high.

There are other variables that might have an effect on the impact of the policy. I check if individuals with a highly educated father have a different intensity of the effect. The idea is that the father's education is a proxy for household income. However, I do not find any significant difference. On the one hand, individuals with an educated father are more likely to have the means to pay the fines. But they are also more likely to have lower fertility preferences. As an interesting addition, I find that individuals whose father is a member of the Communist party are not significantly affected by eligibility. It is plausible that they have already internalized the rule of one child per family and thus their fertility expectations are not affected by a change in eligibility rules.¹⁵

¹⁵It is also possible that they always plan with two children.

1.6.2 Effect of eligibility on fertility outcome

I use the older sample to estimate the effect of actual eligibility on fertility outcome. N_i is measured as having a second child in 2010 or not. The analysis is restricted to married individuals that had one child before 2003. This is to make sure they have enough time to have a second child while conforming with promoted birth intervals and do not have more than two children. Eligibility is measured at the age of 27 for women and 28 for men which is the year before the average age of the birth of the second child.

In specification 1, the indicator variable of having a second child in 2010 or not is regressed on the eligibility status at the age of 27/28. I control for individual characteristics that determine eligibility, as well as birth year and province fixed effects, and the sex of the first child:

$$N_{i \text{ in } 2010} = \tilde{\beta} Z_{ip \text{ age}(27/27)} + X_i + \gamma_t + \kappa_p + \epsilon_{ipt} \quad (1.26)$$

Equation 1.26 can be estimated by a logit regression. However, the underlying assumptions such that $E(\epsilon_i Z_i | X_i) = 0$ are strong. The effect of the characteristics that make someone eligible such as belong to an ethnic minority or having an agricultural household status are assumed to be the same in each province. Also, year fixed effects are assumed to be the same for all provinces and for all population groups. To relax these assumptions, I also run a DID and double DID specification as in the previous section using a OLS regression due to the high number of fixed effects.

The results are displayed in table 1.3. Eligibility for a second child permit at the age of 27/28 has a highly significant effect in all specifications (columns 1 to 3). On average, being eligible increases the likelihood to have a second child by 14 percentage points, based on the DID specification. Effects are similar for men and women (columns 5 and 6).

The official number of children allowed at age 27/28 thus influences real fertility decisions. However, an increase of approximately 14 percentage points implies that there is a share of the population that does not significantly change their fertility outcomes due to the policy. This might be because they want only one child or because they were planning to have two children and to pay the fine.

The coefficient measures only the impact of eligibility based on criteria that are observable at the age of 16. There are several criteria that are based on unobservable and (for the econometrician) not predictable characteristics. These include the first child being disabled and the family being in “poor economic conditions”. Furthermore, eligibility is only measured at 27/28 which is the year before the average age of the birth of a second child. This implies that individuals who become eligible at the age of 28/29 fall into the control group. The same applies to those who become non-eligible at age 27/28 but might have already had a second child or had gotten a second child

Dependent variable: Indicator: Having a second child						
	OLS DID (1)	Logit (2)	OLS Double DID (3)	OLS Double DID (4)	OLS-DID Women (5)	OLS-DID Men (6)
Eligibility at age 27(women)/28(men)	0.143*** (0.0438)	0.110*** (0.0284)	0.152*** (0.0449)	0.144*** (0.0447)	0.153*** (0.0483)	0.139*** (0.0477)
Eligibility Controls	Yes	Yes	Yes	Yes	Yes	Yes
Additional Controls	Yes	Yes	Yes	Yes	Yes	Yes
Local Fertility Rate	No	No	No	Yes	No	No
Province FE	Yes	Yes	Yes	Yes	Yes	Yes
Elig. Controls x Province FE	Yes	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Province FE x Year FE	Yes	No	Yes	Yes	Yes	Yes
Elig. Controls x Time FE	No	No	Yes	Yes	No	No
Observations	8423	8423	8192	8423	4554	3869
R^2	0.420		0.447	0.434	0.457	0.442

Note: Dependent variable = 1 if the individual had a second child. Coefficients are average marginal effects. Sample includes individuals that turned 16 between 1977 and 1992, are married, had their first child before 2003 and have no more than two children. Logit or OLS regressions with standard errors in parenthesis. Standard errors are clustered on the province interacted with an urban dummy which indicated that in 2010, the area is defined as an urban area. Eligibility at age 27/28 is calculated on the basis of the province, eligibility characteristics, sex of the first child and time. Eligibility controls: Household status, only child status, minority status. Additional controls: sex, sex of the first child. Data source: China Family Panel Survey 2010.

Table 1.3: Effect of eligibility status at age 27/28 on the likelihood of having a second child.

permits before the change. Therefore, the coefficient does not measure the complete impact of 2nd child permits and probably underestimates even the effect of the criteria included.

1.6.3 Isolating the effect on *increasers*

To decompose the effect of eligibility on education, I want to split the sample into two groups: those that are expected to be *always-2* and those that are not. I use the subsample of those not eligible in the older sample as a training set for a random forest estimation. It predicts if someone has a second child without being eligible. Only those that turn 16 before 1990 are included so that the training sample and the main sample do not overlap. Motivated by proposition 2 in the theory section, I include potential fertility proxies on top of eligibility characteristics, province indicators and a trend variable. The local, district-level fertility rate is calculated as the average number of children of those between 30 and 35 (when the individual turns 16). I also include the number of siblings and an indicator if the mother has finished at least junior high school. Also included are if the individual grew up in an urban area (independent of hukou status) and if the father is a member of the Communist party.

The random forest grows 1000 classification trees. For each tree 3 predictors are used which are randomly selected from all predictors. The number of predictors is chosen to minimize the out-of-bag error rate. Prediction is made based on the average of those trees. An observation is predicted to be an *always-2* if more than half of the observation that falls into the same category have a second child (majority rule). Figure 1.7 shows the importance of the different predictors. Unsurprisingly, the province and the local fertility rate are the most important predictors. The trend, number of siblings, and the household registration status are important as well.

The random forest is then used to predict if an individual in the older sample and the main sample would be an *always-2*. This variable is then interacted with the eligibility status. Results are displayed in table 1.4. The first two columns (1 and 2) show the effect of eligibility at age 27/28 on the indicator of having a second child. As expected, the take-up rate among those that are not predicted to be *always-2* is significantly higher than among those predicted *always-2*. For those predicted to be *always-2* second child permits have an effect close to zero.

The effect of anticipated eligibility on education is still positive and significant for those that are not predicted to be *always-2* (column 3 and 4). The interaction has a negative sign but is not significant. YET, the effect of eligibility on senior high school enrollment is significantly higher for those that are not predicted to be *always-2* (columns 5 and 6). These results suggest that the positive effect of eligibility on educational investment is driven by those that increase their anticipated fertility as a response to eligibility status.

Sample	Having a 2nd child		Dependent variable:			
			Years of Education		Enrollment into SHS	
	1977-1992 DID (1)	1977-1992 Double DID (2)	1990-2005 DID (3)	1990-2005 Double DID (4)	1990-2005 DID (5)	1990-2005 Double DID (6)
Eligibility at age 27/28	0.180*** (0.0511)	0.183*** (0.0512)				
Eligibility at age 27/28 X Predicted <i>Always-2</i>	-0.125** (0.0535)	-0.129** (0.0545)				
(Anticip.) eligibility at age 16			1.048* (0.537)	1.212*** (0.440)	0.134** (0.0651)	0.120* (0.0652)
Predicted <i>Always-2</i> X (Anticip.) eligibility at age 16			-0.511 (0.384)	-0.475 (0.375)	-0.0914*** (0.0323)	-0.0942*** (0.0309)
Predicted <i>Always-2</i>	0.793*** (0.0176)	0.794*** (0.0180)	-0.836** (0.347)	-0.929*** (0.329)	-0.0485 (0.0304)	-0.0448 (0.0298)
Observations	8104	8104	7528	7528	7528	7528
R^2	0.631	0.633	0.472	0.466	0.423	0.428

Note: Includes eligibility characteristics, province and time fixed effects, province X time fixed effects, eligibility characteristics X province fixed effects, female dummy. Anticipated eligibility at age 16 is calculated on the basis of the province, eligibility characteristics and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. Column 1 and 2 include sex of first child, column 2, 4 and 6 include eligibility characteristics X time fixed effects. Standard OLS regression with standard errors clustered on provinces-times-urban-area levels in parenthesis. Significance levels: * 0.10; ** 0.05; *** 0.01. Data source: China Family Panel Survey 2010.

Table 1.4: Decomposing the effect of eligibility

1.6.4 Effect of eligibility and the sex ratio

The theoretical model predicts heterogeneity in the cost of raising children, in the change in labour supply and in the sex ratio. I investigate the latter, as the sample sex ratio is observable and plausibly exogenous on the individual level.¹⁶ I use the 2000 Census to calculate the share of women for each birth year.¹⁷ For the 1977-1992 sample, the average share of women is 0.495, for the 1990-2005 sample the average share is 0.488. The geographical variation of the 1990-2005 sample is displayed in figure 1.6.

At first, I check if the sex ratio has an effect on the take-up rate of second child permits (see table 1.11, column 1). I find that the interaction between eligibility and sex ratio is insignificant. I still include the sex ratio in the random forest prediction for *always-2*. It is less important than the local fertility rate, province fixed effects or the household status, but as important as the number of siblings and the trend variable (see figure 1.8).

I find that the sex ratio has a significant effect on the impact of second child permits. The interaction term between sex ratio and anticipated eligibility status at age 16 is negative and significant (see table 1.11 column 2). Thus, the fewer women compared to men are in the cohort, the stronger the effect of second child permits. This is indeed driven by the more frequent sex (men). In columns 3 and 4 where I look at men and women separately, the coefficient is significant (and also larger in size) for men. If the effect is driven by *always-2* or not is verified in columns 5 to 7. Splitting up the sample further decreases power. However, the size of the coefficient for men that are not predicted *always-2* stays the same (though it is not significant) in column 7 (p-value = 0.22).

1.7 Discussion: Why would those that expect more children educate more?

The theoretical model provides different mechanisms of how anticipated fertility can affect educational investment. It can have an effect on family consumption, the labour supply and the incentives to find a partner. In this section, I provide descriptive statistics to discuss the intertemporal consumption smoothing and labour supply channels. Looking at expenditure data and parental labour supply helps to put the results into context.

The cost of raising children not only depends on how much money is spent on the need of the children but also on the available income at that moment. If parents also have to financially support their own parents while raising their children, it makes the financial burden of children more

¹⁶Expenditures for the second child are not observed and endogenous, as are labour market conditions.

¹⁷The census, having more observations than the CFPS, gives a cleaner measure.

relevant. It is also more important for the grandparents that have a strong say in the educational decision making. Therefore, I also look at the financial and living situation of the elderly.

1.7.1 Cost of raising children

The fertility literature in developing countries has treated children as an investment for parents that pay off quickly because children are productive from an early stage on. The idea of children as productive household members relies on children being able to work from an early stage on. In developed countries, children are often seen as expensive. Children have to be financially supported until the end of the educational stage. Furthermore, retirement schemes and health insurance mitigate the need to rely on children during sickness or old age.

In China, education is mandatory until age 16 and children are not allowed to work. If families abide by these rules, they have additional expenses for at least 16 years per child. Indeed, gross enrolment rates for primary school have constantly been above 100% for the past years. At the secondary school level, they have increased from around 70% to over 90% between 2007 and 2015 and for tertiary education from 20% to 45% (UNICEF). At the same time, child labour is not prevalent. Tang et al. (2016) estimate that in 2010, only about 7.74% of children aged from 10 to 15 were working. Most of them combined education with their economic activity.

Expenditure on the education of their children is an important cost for parents in many countries. Figure 1.9 in the appendix shows total expenditure on education as a share of family income for families with one child. Though educational expenditure is low in the first four years, 10 to 20% of family income is spent on education afterwards. Figure 1.10 shows the different expenditure items as a share of family expenditure. Expenditure increases at the age of 5, probably driven by tuition fees for pre-school. Tuition drops at age 6 or 7 when children start primary school. At the same time families spend more on extracurricular activities and home tutoring. Tuition becomes more important again at the age of 17 when students enter senior high school, then college or university.

Figure 1.11 in the appendix contrasts the educational expenditures of a family with two children with a family with only one child. To compare these families, I use a nearest-neighbour matching technique.¹⁸ Families with two children pay more for education than families with one child starting when the second child is around 5 years old. When the second child is around 15/16 years old, educational expenditure accounts for more than 20% of family income.

¹⁸First, the likelihood of having a second child is regressed on age, agricultural household status, educational level and age at the birth of the first child and the propensity score is predicted. Each individual is then matched with their nearest neighbour (of the same sex) without replacement based on the propensity score. The matched control individual “adopts” the birth year of the second child from their neighbour.

This descriptive result is replicated in a regression analysis displayed in table 1.12 in the appendix. Using standard OLS regression, I find that families with a second child spend twice as much on education than families with only one child. I instrument the number of children with eligibility at age 27/28 to address the endogeneity issue. The result stays the same (column 3). The IV approach is only suggestive. Eligibility potentially affects other variables that influence educational expenditure that I cannot control for, such as savings.

A recent report estimates that Chinese parents pay around \$42,892 for their child's education (HSBC (2017)). Though this is less than in the USA (\$58,464) where college attendance is higher, it is more than in other countries such as the UK (\$24,862), Canada (\$22,602) or France (\$16,708) (HSBC (2017)). Furthermore, the report revealed that 93% of Chinese parents have paid for a private tutor. Chinese parents seemed to have expected these expenses. Out of the countries that were surveyed, China had the highest rate of parents (around 55%) indicating that they can fund their child's education through savings or investment.

The suggestive evidence and the additional literature indicate that raising children in China is indeed costly. In particular, families spend an important share of their income on the education of their children.

1.7.2 Old-age support and the double burden of parents

The effect of anticipated fertility also depends on the importance of children's financial support to parents. It makes children more costly when investment in children and financial support for the grand-parents are paid at the same time. Children are still an important source for old age support in China. According to the CFPS, in 2010, 45% of individuals over 60 live with at least one of their children, 69% of those with a son and 31% with a daughter. For families with an agricultural household status and at least one person older than 60, social security/pensions transfers accumulate to on average 2068 Yuan (approx. 218€)¹⁹, making up less than 10% of family income. For families with a non-agricultural household status, it accumulates to on average 20353 Yuan (2137€, around 47% of family income). 33% of individuals over 60 states having received economic support from at least one of their children.

Families with both spouses born between 1964 and 1979 give 4% of their annual income to relatives. They increase family member support by 1 percentage point for each grandparent that does not live with them (see table 1.14 in the appendix). Using OLS, there is no significant difference in the amount given to relatives between families with one child and two children. Once having a second child is instrumented (column 3), I find that families with one child give around 3 per cent of family income to a grandparent who does not live with them. This is not the case for families with two children. The expenses for the second child seem to impact the ability to financially support

¹⁹Based on the conversion rate of 0.105, the approximate average exchange rate at the beginning of 2010.

grandparents.

1.7.3 Parental labour supply and earnings

Do parents work more or less when they have a second child? The CFPS includes data on how much the individual worked on average per month in 2010. Figure 1.12 in the appendix shows the average log working hours of men and women as a function of how many years have passed since the birth of the second child. It is displayed for parents that have exactly two children and a matched control group with one child. If at all, men work significantly less only at the year of birth of the second child but not afterwards.²⁰ Women work significantly less in the first four years following the birth of the second child. Then they work the same hours or potentially even more from year 5 onward. Figure 1.13 in the appendix displays parents' monthly income. In line with the previous results, men earn less in the year of birth and there is no evidence for a lasting effect. Women's earnings are less in the first four years after the birth of the second child. Long-run earning patterns are noisy: there is no clear trend for either higher or lower earnings.

Again, the descriptive results are replicated in a regression analysis (see table 1.13 in the appendix). Using standard OLS, I find that having a second child is negatively associated with working hours for men and women in 2010 (columns 1, 2, 4 and 5). When having a second child is instrumented, the coefficient for women is no longer significant but still negative (column 3). The effect of having a second child for men is insignificant with a positive sign (column 6).

Guo et al. (2017) and He and Zhu (2016) address the endogeneity between childbirth and labour market conditions by using twinning as a natural experiment. Guo et al. (2017) do not find evidence on the negative effects of fertility on parental labour supply. He and Zhu (2016) find a small negative effect on women's labour force participation in the 1990s and an insignificant effect in 2000s. Angrist and Evans (1996), Jacobsen et al. (1999) and Lundborg et al. (2017) show for other countries that one more child has no or only a small effect on labor market earnings of women.

Overall, the descriptive statistics, suggestive IV results and the related literature suggests that in the Chinese context having a second child does not have a long-run effect on parental labour supply and earnings. Therefore, the positive effect of an increase in anticipated fertility on education is not surprising.

1.8 Conclusion

In this chapter, I use a novel empirical approach to address the question of how anticipated fertility affects educational decisions. This question has not been in the focus of the fertility literature. I use the One-Child Policy in China and the existence of second child permits for a subset of

²⁰There are other years where they work significantly less at 10% but the difference is not systematic.

individuals. I find that individuals invest more in education when they fulfil an eligibility criterion for a second child permit. Treatment heterogeneity analysis suggests that the effect is driven by compliers: those who increase their anticipated fertility once eligible.

In the theory section, I sketch economic channels of how anticipated fertility can positively influence educational investment. Because children are expensive one wants to ensure earning sufficient income in the future. The positive effect also depends on the relationship between lifetime returns to education and fertility outcomes. If men and women can stay in or re-enter the labour market without loss of their skills one does not expect a negative effect. The policy implication is thus to ensure the availability of flexible childcare, to fight potential discrimination of parents in the labour market and to support parents who want to re-enter the labour market.

There are other channels of how anticipated fertility can affect education. For one, individuals who plan to have more than one child in the future might also plan to take over important childcare tasks. They invest more in education to be well prepared for it. Also, one parent might want to increase their bargaining power when they count on having to secure sufficient resources for more than one child. Since bargaining power within the household and education are supposedly positively correlated, increasing education can be a way to increase bargaining power. There is also a potential psychological effect. For example, being allowed to have two children in a society where children are seen as essential can give one a more positive attitude towards the future and more motivation at school.

Similarly, there are many other variables that anticipated fertility could have an effect on. Do men and women who anticipate several children choose certain types of jobs or avoid certain sectors? Do saving patterns change? Do they look for different types of partners? These questions are important to ensure that individuals have the same opportunities independent of their desire to have children. They can also help understand saving, investment and marriage patterns.

China is one specific social and economic environment that has been perturbed by strict policies. Comparing individuals who plan to have two children instead of one (intensive margin) is not the same as comparing individuals who do not plan to have any children with those that do (extensive margin). Studies from developed countries find that having one more child has at most a weak effect on the mother's long-run ([Angrist and Evans \(1996\)](#), [Jacobsen et al. \(1999\)](#), [Lundborg et al. \(2017\)](#)). However, the first child has a significant effect. This suggests that the effect of having a child compared to not having any on education can also be different.

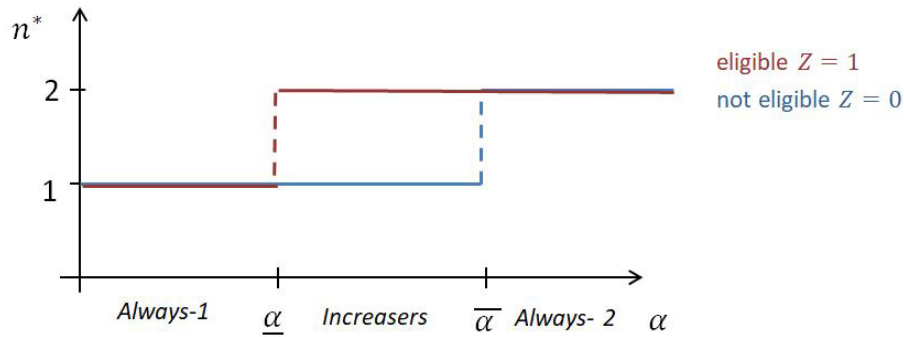
China announced a Two-children policy in 2015, which allows everyone to apply for a second child permit. There is an ongoing discussion about the effect of the OCP on fertility rates and thus the effect of the new policy. The results of this chapter suggest that the Two-children policy

has an effect on fertility, but it is expected to be small. The policy is also predicted to increase educational investment.

The findings of this chapter are a result of the Chinese social and economic environment. However, low fertility is a concern in many medium and high-income countries. This finding is a positive sign for policymakers who want to promote fertility and education at the same time. Identification relies on a setting with an exogenous, predictable and important variation in the cost of a child. This will make replication in other countries difficult. Panel studies can be used which ask for fertility goals and when they can be coupled with changes in fertility policies. A dynamic and structural framework can be applied in this context. This chapter is one of the first empirical papers that address the identification issue of the relationship between anticipated fertility and educational investment. Further research will hopefully follow.

1.9 Appendix

1.9.1 Figures



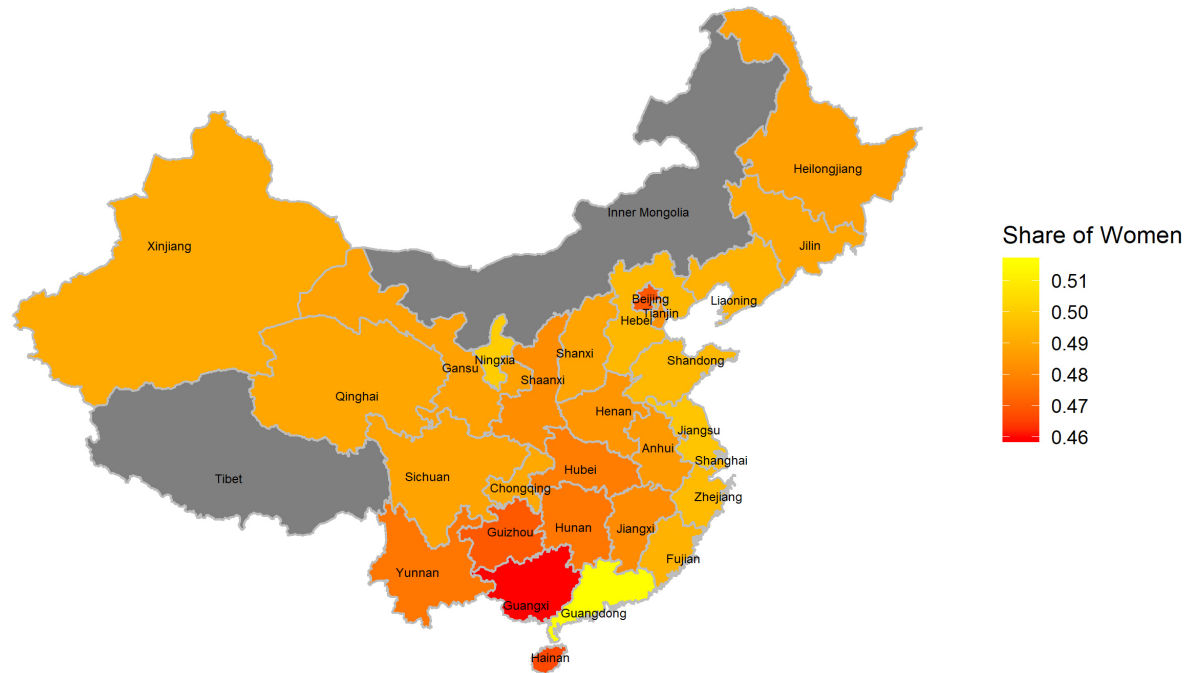
Note: Optimal number of children n^* on the y axis and fertility preferences α on the x axis. The red (blue) lines demonstrate the optimal number of children as a function of fertility preferences when the individual is eligible (not eligible). The thresholds $\bar{\alpha}$ and $\underline{\alpha}$ are determined by the FOC.

Figure 1.4: The effect of fertility preferences on the number of children.



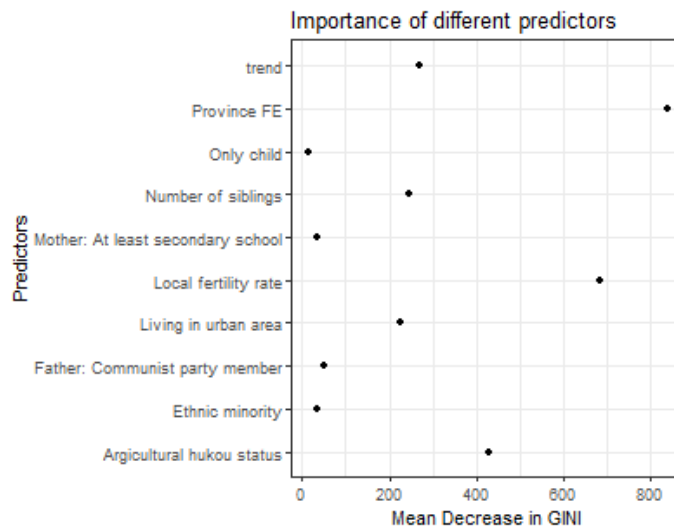
Note: Effect size in percentage points of anticipated fertility at the different age thresholds 14, 15, 16, 17 and 18. Based on DID specification with province level trends, standards errors clustered on province times urban level. 90% Confidence intervals.

Figure 1.5: Effect size of anticipated eligibility on likelihood of senior high school enrolment at different age thresholds.



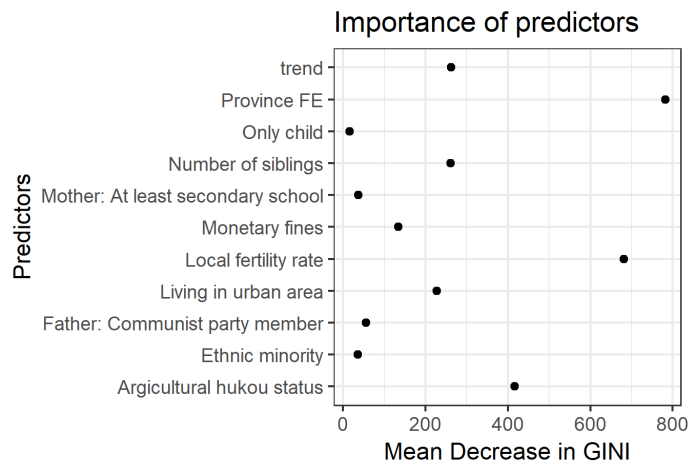
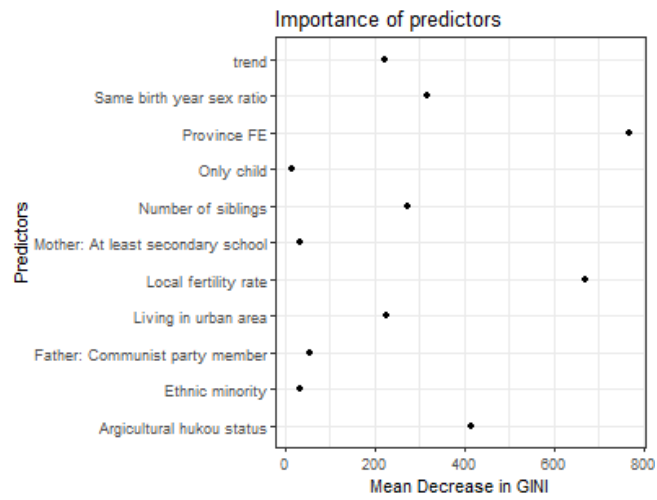
Note: Sex ratio based on the 2000 Census.

Figure 1.6: Average sex ratio for those who turned 16 between 1990 and 2000.



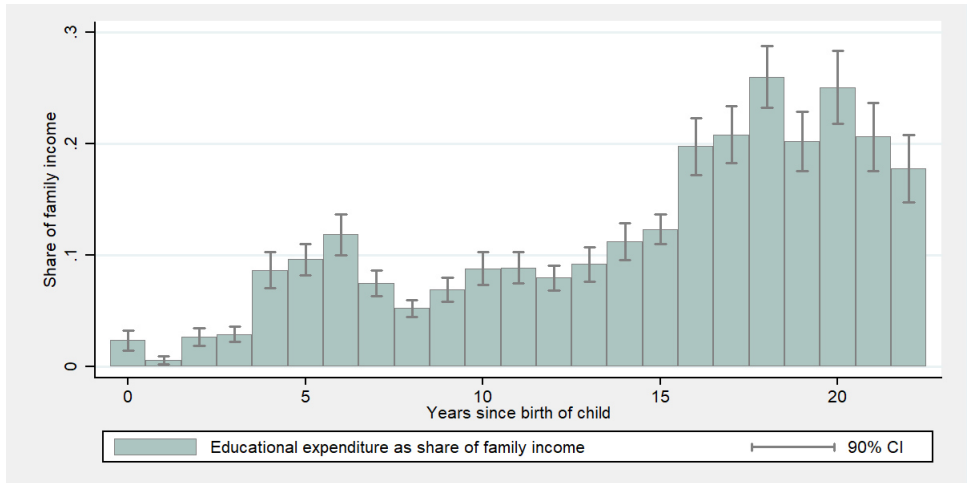
Note: Displays the importance of predictors in the random forest estimation. Importance is measured as how much the predictor decreases the GINI coefficient within each node. If all observations at one node are either 0 or 1, GINI is 0. If the observations are mixed, the GINI coefficient is high. The GINI is a measure of node purity.

Figure 1.7: Importance of predictors in random forrest estimation



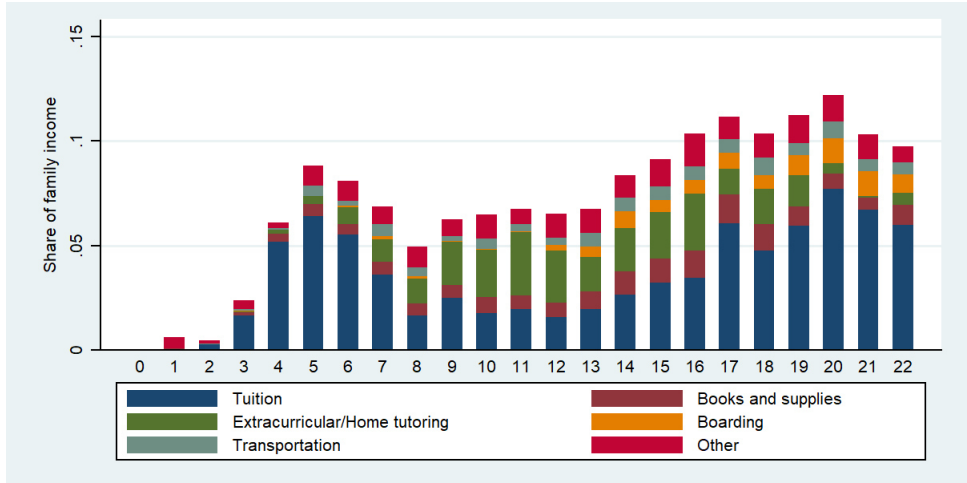
Note: Displays the importance of predictors in the random forest estimation. See figure 1.7

Figure 1.8: Importance of predictors including the cohort sex ratio (from 2000 census) on the top and monetary fines on the bottom



Note: Share of family income allocated to education for families that have an only child. Data is based on the family questionnaire and the question of how much was spend on education last year. Source: CFPS 2010.

Figure 1.9: Total educational expenditure as share of family expenditure




Note: Share of family expenditure allocated to education for families that have an only child. Data is based on the child questionnaire until the age of 15 and the adult questionnaire from age 16 onward. Includes families with an only child where the child lives with at least one parent. Excludes families that do not pays all the educational cost (28 observations). The question is asked for each educational item. Source: CFPS 2010.

Figure 1.10: Different items of educational expenditure as share of family expenditure.



Note: Share of family expenditure allocated to education comparing those with two children with a matched comparison group. The comparison group is matched 1 to 1 without replacement based on the propensity score taking into account age, education, household registration status and the age at the birth of the first child. The matched control observation adopts the year of birth of the second child. Data is based on the family questionnaire of the CFPS 2010.

Figure 1.11: Expenditure on education for families with two children and a matched control group with one child



Note: Log of monthly working hours for men and women with 90% confidence intervals. Blue is the average of those who have two children. Red is a matched control group with one child. The comparison group is matched 1 to 1 without replacement based on the propensity scores taking into account age, education, household registration status and the age at the birth of the first child. The matched control observation adopts the year of birth of the second child. Data source: CFPS 2010.

Figure 1.12: Log monthly working hours in 2010 according to the year of birth of the second child



Log of monthly income for men and women with 90% confidence intervals. Blue is the average of those who have two children. Red is a matched control group with one child. The comparison group is matched 1 to 1 without replacement based on the propensity scores taking into account age, education, household registration status and the age at the birth of the first child. The matched control observation adopts the year of birth of the second child. Data source: CFPS 2010.

Figure 1.13: Log monthly income in 2010 according to the year of birth of the second child

1.9.2 Tables

	Hebei	Chongqing	Hubei	Zhejiang	Jiangsu
Family only has girl (rural area)	1989	1997	1987	1995	2002 ²¹
Ethnic minorities	1982	2002	2002	1990	-
Spouses are only child	1982	1997	2002	1989	1990

Note: Provinces have several other eligibility criteria such as for couples who had their first child outside of China, remarried couples, couples with a disabled first child etc. that I do not regard. Based on [Scharping \(2013\)](#) and family planning documents.

Table 1.5: Example of when provinces formalized eligibility criteria.

²¹Only if husband does not have a brother.

²²Truncated at 7

²³Taking into account the sex of the first child.

Variable	Mean	Std. Dev.	Min.	Max.
Main Sample (turning 16 between 1990 and 2005)				
Years of education completed (in 2010)	8.681	4.631	0	22
Female	0.534	0.499	0	1
Han ethnicity	0.893	0.309	0	1
Rural/agricultural household registration status	0.689	0.463	0	1
Year born	1981.291	4.738	1974	1989
Nb children allowed: 2 (at age 16)	0.275	0.446	0	1
Nb children allowed: 1.5 (at age 16)	0.376	0.484	0	1
N	7840			
Older Sample: (turning 16 between 1977 and 1992)				
Years of education completed (in 2010)	6.527	4.664	0	22
Female	0.529	0.499	0	1
Han ethnicity	0.915	0.278	0	1
Rural/agricultural household registration status	0.732	0.443	0	1
Year born	1968.208	4.346	1961	1976
Nb children allowed: 2 (at age 16)	0.137	0.343	0	1
Nb children allowed: 1.5 (at age 16)	0.11	0.309	0	1
Number of children ²²	1.673	0.821	0	7
Allowed to have 2nd child ²³	0.36	0.48	0	1
N	11626			

Note: Data source China Family Panel Study 2010. Only those with information on the main variables (education, year of birth etc).

Table 1.6: Summary statistics

	Dependent variable: Indicator of enrollment/completion in (at least)							
	Junior High School Completion		Senior High School Enrollment		College or University Enrollment		University Enrollment	
	DID (1)	Double DID (2)	DID (3)	Double DID (4)	DID (5)	Double DID (6)	DID (7)	Double DID (8)
Anticip. eligibility at age 16	0.0272 (0.0368)	0.0254 (0.0376)	0.106* (0.0601)	0.0898 ⁺ (0.0610)	0.0967** (0.0399)	0.0981** (0.0398)	0.0611** (0.0280)	0.0658** (0.0291)
Female dummy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Eligibility controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Eligib. controls x Province FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE x Province FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Eligibility controls x Year FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	7840	7840	7840	7840	7840	7840	7840	7840
R^2	0.312	0.318	0.418	0.424	0.324	0.340	0.167	0.184

Note: Sample includes individuals that turned 16 between 1990 and 2005. Dependent variable is the likelihood of completing junior high school (columns 1 and 2), enrolling into senior high school (columns 3 and 4) enrolling into (technical) college or university (columns 4 and 5) and university (columns 7 and 8). OLS regressions with standard errors in parenthesis. Standard errors are clustered on the province times urban level. Anticipated eligibility at age 16 is calculated on the basis of the province, eligibility characteristics and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. Significance levels: * 0.10; ** 0.05; *** 0.01. Data source: CFPS 2010.

Table 1.7: The effect of the eligibility on the likelihood of completing a degree.

	Dependent variable: Years of Education	
	(1)	(2)
Anticipated eligibility at age 16	-0.112 (0.592)	-0.791 (0.616)
Female dummy	Yes	Yes
Eligibility controls	Yes	Yes
Province FE	Yes	Yes
Eligibility controls x Province FE	Yes	Yes
Year FE	Yes	Yes
Year FE x Province FE	Yes	Yes
Eligibility controls x Year FE	No	Yes
Observations	6705	6903
R^2	0.332	0.321

Note: Sample includes individuals that turned 16 between 1970 and 1980. Placebo eligibility is calculated as if the individuals were born 20 years later (between 1990 and 2000). Dependent variable is the years of education the individual completed. Standard OLS regression with robust standard errors in parenthesis. Significance levels: * 0.10; ** 0.05; *** 0.01 Data source: China Family Panel Survey 2010.

Table 1.8: Placebo test.

	Dependent variable: Years of Education			
	(1)	(2)	(3)	(4)
(Anticip.) Eligibility at age 16 = 0.5	0.258 (0.263)			
(Anticip.) Eligibility at age 16 = 1	0.961** (0.382)			
(Anticip.) Eligibility at age 16		0.815** (0.381)	1.110*** (0.385)	1.337* (0.693)
Agri. household status X (Anticip.) Eligibility age 16		0.277 (0.318)		
Only Child X (Anticip.) Eligibility (Anticip.) age 16			-1.469** (0.723)	
Han ethnicity X (Anticip.) Eligibility age 16				-0.545 (0.690)
Female dummy	Yes	Yes	Yes	Yes
Eligibility controls	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes
Eligibility controls x Province FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE x Province FE	Yes	Yes	Yes	Yes
Eligibility controls x Year FE	Yes	Yes	Yes	Yes
Observations	7840	7840	7840	7840
R^2	0.467	0.467	0.468	0.467

Note: Sample includes individuals that turned 16 between 1990 and 2000. Dependent variable is the years of education the individual completed. Standard OLS regression with robust standard errors in parenthesis. Double DID specification. Significance levels: * 0.10; ** 0.05; *** 0.01 Data source: China Family Panel Survey 2010.

Table 1.9: Effect of anticipated eligibility on the years of education - Robustness Checks.

	Dependent variable: Years of education	
	Double DID (1)	Double DID (2)
Anticip. Eligibility at 16	1.129** (0.454)	1.099** (0.532)
Criteria dependent on spouse	-0.295 (0.312)	-0.283 (0.307)
Double DID controls	Yes	Yes
Observations	7840	7840
R^2	0.462	0.467

Note: Double difference controls: year fixed effects, province fixed effect, province times year fixed effects, eligibility controls, eligibility controls times year fixed effects and eligibility times province fixed effects. Sample includes individuals that turned 16 between 1990 and 2005. The dependent variable is the years of education the individual completed in 2010. Standard OLS regression with standard errors clustered on provinces-times-urban-area levels in parenthesis. Significance levels: * 0.10; ** 0.05; *** 0.01. Anticipated eligibility at age 16 is calculated on the basis of the province, eligibility characteristics and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. Data source: CFPS 2010.

Table 1.10: Spouse dependence in the eligibility criteria.

Sample	Dependent variable:						
	Having 2nd child	Enrollment in Senior High School					
	1977-1992 all (1)	1977-1992 all (2)	1990-2005 Women (3)	1990-2005 Men (4)	1990-2005 all (5)	1990-2005 Women (6)	1990-2005 Men (7)
Eligibility at age 27/28	0.124*** (0.0460)						
Eligibility at age 27/28 X Sex ratio (std)	0.0369 (0.0294)						
(Anticip.) eligibility at age 16		0.112* (0.0643)	0.0793 (0.0776)	0.127* (0.0699)	0.136** (0.0619)	0.107 (0.0745)	0.133* (0.0718)
(Anticip.) eligibility at age 16 X Sex ratio (std)		-0.0219* (0.0124)	-0.0105 (0.0176)	-0.0355*** (0.0127)	-0.0177 (0.0207)	0.00260 (0.0279)	-0.0367 (0.0296)
Predicted <i>Always-2</i> X (Anticip.) eligibility at age 16					-0.0787* (0.0395)	-0.0755 (0.0508)	-0.0723 (0.0548)
Predicted <i>Always-2</i> X Sex ratio (std)					-0.0142 (0.0167)	0.00130 (0.0262)	-0.0333* (0.0180)
Predicted <i>Always-2</i> X (Anticip.) eligibility at age 16 X Sex ratio (std)					0.00270 (0.0264)	-0.0225 (0.0366)	0.0206 (0.0371)
Predicted <i>Always-2</i>					-0.0750** (0.0300)	-0.0741*** (0.0264)	-0.0645 (0.0408)
Observations	8423	7840	4184	3656	7528	4015	3513
R^2	0.421	0.419	0.480	0.428	0.424	0.485	0.435

Note: All regressions are OLS Difference-in-Difference specification with province level times fixed effects, control for eligibility characteristics times province fixed effects and gender. Column 1 controls for sex of the first child, includes only those that are married and had their first child before 2003. Anticipated eligibility at age 16 is calculated on the basis of the province, eligibility characteristics and time. Eligibility at age 27/28 is calculated on the basis of province, eligibility characteristics, sex of the first child and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. Cohort sex ratio is calculated as the sex ratio of those born up to 2 years before and after the individual in the same province, based on the 2000 Census. Sex ratio is standardized. Clustered standard errors on province X urban level in parenthesis. Significance levels: * 0.10; ** 0.05; *** 0.01 Data source: CFPS 2010.

Table 1.11: Sex ratio and eligibility.

Sample (years they turn 16)	Dependent variable: Log(expenses on education)		
	OLS	OLS	IV
	1980 -1995	1980 -1995	1980 -1995
	(1)	(2)	(3)
Indicator: Having a second child	1.020*** (0.0876)	1.057*** (0.0892)	0.975** (0.405)
Eligibility Controls	Yes	Yes	Yes
Additional Controls	Yes	Yes	Yes
Province FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Province FE X Eligibility controls	No	Yes	No
Province FE x Year FE	No	Yes	No
Year FE x Eligibility controls	No	Yes	No
Observations	8471	8471	8471
R^2	0.112	0.185	0.112

Note: Sample includes individuals that turned 16 between 1980 and 1995. Robust standard errors in parenthesis. Additional controls: Sex, educational level, educational level of spouse, age at marriage, age at first birth, house ownership, number of adults in the household, if partner has a job, sex of the first child. In column 3: First stage run by logit regression, instrument: eligibility at age 27/28. T-stat: 5.57; F-stat: 69.18). Data source: China Family Panel Survey 2010.

Table 1.12: Effect of the second child policy on educational expenditure.

	Dependent variable:					
	Women: Log working hours			Men: Log working hours		
	OLS	OLS	IV	OLS	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Indicator: Having a 2nd child	-0.259*** (0.0973)	-0.258** (0.106)	-0.139 (0.420)	-0.237** (0.0950)	-0.183* (0.102)	0.333 (0.289)
Eligibility Controls	Yes	Yes	Yes	Yes	Yes	Yes
Additional Controls	Yes	Yes	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Province FE X Eligibility controls	No	Yes	No	No	Yes	No
Province FE x Year FE	No	Yes	No	No	Yes	No
Year FE x Eligibility controls	No	Yes	No	No	Yes	No
Observations	4424	4424	4424	3892	3892	8316
R^2	0.080	0.169	0.080	0.096	0.211	0.064

Note: Sample includes individuals that turned 16 between 1980 and 1995. Robust standard errors in parenthesis. Additional controls: Sex, educational level, educational level of spouse, age at marriage, age at first birth, house ownership, number of adults in the household, if partner has a job, sex of the first child. In column 3 and 6: First stage run by logit regression, instrument: eligibility at age27/28. T-stat: 3.96/4.03; Chi-stat: 15.65/16.28). Data source: China Family Panel Survey 2010.

Table 1.13: Effect of a second child on parental working hours.

Dependent variable: Share of family income spend to support relatives			
Sample (years they turn 16)	OLS 1980 -1995 (1)	OLS 1980 -1995 (2)	IV 1980 -1995 (3)
Indicator: Having 2nd child	-0.00313 (0.00636)	0.000427 (0.0240)	0.0218 (0.0221)
Nb of individual's grandparents living somewhere else	0.0105*** (0.00353)	0.0113*** (0.00373)	0.0321*** (0.0116)
Indicator: Having 2nd child X Nb of individual's grandparents living somewhere else	-0.000790 (0.00514)	-0.00209 (0.00497)	-0.0450** (0.0204)
Eligibility Controls	Yes	Yes	Yes
Additional Controls	Yes	Yes	Yes
Province FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Province FE X Eligibility controls	No	Yes	No
Province FE x Year FE	No	Yes	No
Year FE x Eligibility controls	No	Yes	No
Observations	8104	8104	8104
R^2	0.028	0.076	0.021

Note: Sample includes the individuals that turned 16 between 1980 and 1995. Robust standard errors in parenthesis. Additional controls: Sex, educational level, educational level of spouse, age at marriage, age at first birth, house ownership, number of adults in the household, if partner has a job, sex of the first child. In column 3: First stage run by logit regression, 1. instrument: eligibility at age 27/28, first stage run by logit regression: T-stat: 5.62; Chi2-stat: 31.64; 2.instrument: Nb of individual's grandparents living somewhere else x eligibility at age 27/28. First stage run by OLS: T-stat: 8.42; F-stat: 70.91. Data source: China Family Panel Survey 2010.

Table 1.14: Effect of the second child policy on household balance.

1.9.3 Model appendix

1.9.3.1 The effect of another child on optimal education

Child care generally is time-intensive and parents might forgo earnings because they have to reduce working hours. Let us use the example income-generation function from the model section: $Y(I, n) = (T - \delta n)\rho P(I)$, where T is the maximum working hours and for every child the parents has to spend δ hours on child care. Each working hour is remunerated with $\rho P(I)$ where $P(I)$ is the productivity, which is a function of human capital, and ρ is a scaling parameter which can be interpreted as technology or labour market conditions. So parents that have one more child, have a reduction in income:

$$\frac{\partial Y(I, n)}{\partial n} = -\delta\rho P(I) \quad (1.27)$$

We also define $f(n)$ as $\mu F(n)$ where μ captures difference in the cost, and $F(n)$ is an increasing, concave function in n . In this section, we ignore savings.

We can see that given the example income-generation function, the cost of having to care for a child is a function of human capital. Looking at the optimal educational level (disregarding penalties for the second child):

$$u'(y(I^*, n) - \mu F(n)) \left[\frac{\partial y(I^*, n)}{\partial I^*} \right] = \frac{\eta}{\delta} u'(Y - \eta I^*) \quad (1.28)$$

We now look at how the optimal educational level changes with n :

$$\underbrace{u''[y(I^*, n) - \mu F(n)]}_{(-)} \left[\underbrace{\frac{\partial y(I^*, n)}{\partial n} - \mu f'(n)}_{(-)} \right] \underbrace{\frac{\partial y(I^*, n)}{\partial I^*}}_{(+)} + \underbrace{u'(y(I^*, n) - \mu F(n))}_{(+)} \underbrace{\frac{\partial^2 y(I^*, n)}{\partial I^* \partial n}}_{(-)} \quad (1.29)$$

The first part is positive: $u''()$ is negative; $\frac{\partial y(I^*, n)}{\partial n}$ is negative, as well as $-\mu f'(n)$ while $\frac{\partial y(I^*, n)}{\partial I^*}$ is positive. This is the consumption smoothing effect: The cost of raising children directly though $f(n)$ as well as indirectly through the forgone income $\frac{\partial y(I^*, n)}{\partial n}$ decrease income available for consumption and thereby increase the marginal utility of additional earning.

The sign of the second part depends on the sign of $\frac{\partial^2 y(I^*, n)}{\partial I^* \partial n}$. Using our income generation function from before:

$$\frac{\partial^2 y(I^*, n)}{\partial I^* \partial n} = -\delta P'(I) \quad (1.30)$$

Since hourly wage/productivity is increasing in human capital, returns to education decrease in the number of children. The less time the individual can spend working, the less payoff the individual received from investment in education. This is the labour supply effect: reducing working hours due to children decreases the payoffs of educational investment.

The overall effect is positive when the consumption smoothing effect is larger than the labour supply effect:

$$u''(y(I^*, n) - \mu F(n)) \left[\frac{\partial y(I^*, n)}{\partial n} - \mu F'(n) \right] \frac{\partial y(I^*, n)}{\partial I^*} > -u'[y(I^*, n) - \mu F(n)] \frac{\partial^2 y(I^*, n)}{\partial I^* \partial n} \quad (1.31)$$

If we assume constant absolute risk aversion (CARA) with a the coefficient of absolute risk aversion, we get:

$$a \frac{\partial y(I^*, n)}{\partial I^*} \left[\mu F'(n) - \frac{\partial y(I^*, n)}{\partial n} \right] > -\frac{\partial^2 y(I^*, n)}{\partial I^* \partial n} \quad (1.32)$$

1.9.3.2 How does the overall effect change with the cost of raising children μ ?

We first take the total derivative of the left hand side of equation 1.31 (consumption smoothing effect) in terms of the cost parameter μ :

$$\underbrace{u'''(y(I^*, n) - \mu F(n))}_{(+)} \underbrace{\frac{\partial y(I^*, n)}{\partial I^*}}_{(+)} \underbrace{\left[\mu F'(n) - \frac{\partial y(I^*, n)}{\partial n} \mu F(n) \right]}_{+} - \underbrace{u''(y(I^*, n) - \mu F(n))}_{(-)} \underbrace{\frac{\partial y(I^*, n)}{\partial I^*}}_{(+)} \underbrace{F'(n)}_{(+)} \quad (1.33)$$

Under the assumption of prudence ($u'''(\cdot) > 0$), higher costs of raising children strengthens the consumption smoothing channel. This implies that if there is no labour supply effect of children, i.e. $\frac{\partial^2 y(I^*, n)}{\partial I^* \partial n} = 0$ and $\frac{\partial y(I^*, n)}{\partial n} = 0$, then the effect of anticipated fertility on education increases with the cost of raising children (Proposition 1, part 1).

We now take the total derivative of the right hand side of equation 1.31 (labour supply effect) in terms of μ :

$$\underbrace{u''[y(I^*, n) - \mu F(n)]}_{(-)} \underbrace{\frac{\partial^2 y(I^*, n)}{\partial I^* \partial n}}_{(-)} \underbrace{\mu F(n)}_{(+)} \quad (1.34)$$

We can see that the cost of raising children also strengthens also the labour supply channel.

If we assume CARA, as in equation (1.32), the cost of raising children increases the left-hand side, but has no effect of the right-hand side, such that the overall effect is:

$$a \frac{\partial y(I^*, n)}{\partial I^*} F'(n) \quad (1.35)$$

of which all parts are positive.

1.9.3.3 How does the overall effect change with the loss in working time δ ?

In order to investigate the effect of a working time reduction, we assume for simplicity that $f(n) = 0$:

$$u''(y(I^*, n)) \frac{\partial y(I^*, n)}{\partial n} \frac{\partial y(I^*, n)}{\partial I^*} + u'(y(I^*, n)) \frac{\partial^2 y(I^*, n)}{\partial I^* \partial n} \quad (1.36)$$

Using our example income generation function:

$$\underbrace{u''((T - \delta n)P(I^*))(-\delta P(I^*))}_{\text{Income smoothing channel}} \underbrace{(T - \delta n)P'(I^*) - u'((T - \delta n)P(I^*))(\delta P'(I^*))}_{\text{Labour supply channel}} \quad (1.37)$$

Taking total derivatives of the labour supply channel in terms of δ :

$$u'(c^2)P'(I) + u''(c^2)\delta nP'(I)(-nP(I)) \quad (1.38)$$

of which both terms are positive. The labour supply effect is increasing in the loss of working time.

Taking total derivatives of the income smoothing channel in terms of δ :

$$\underbrace{u''(c^2)}_{(-)} \underbrace{P(I)P'(I)}_{(+)} \underbrace{[2\delta n - T]}_{?} + \underbrace{u'''(c^2)}_{(+)} \underbrace{nP(I)\delta P(I)(T - \delta n)P'(I)}_{(+)} \quad (1.39)$$

While the second term is positive under the assumption of prudence, the sign of the first term depends on the sign of $2\delta n - T$. If $T > 2\delta n$, the loss in working hours increases the income smoothing channel (while the reverse is not necessarily true).

We again use at the case of constant absolute risk aversion. Rewrite equation 1.32 without cost of raising a child and using the income generation function:

$$a(T - \delta n)P'(I)\delta P(I) > \delta P'(I) \quad (1.40)$$

which reduces to:

$$a(T - \delta n)P(I) > 1 \quad (1.41)$$

As an increase in δ decreases the left hand side, it decreases the effect of anticipated fertility on education.

1.9.3.4 How does the overall effect change with income?

Suppose the income generation function is $Y(I, n) = (T - \delta n)\tau P(I)$ where τ captures some technology factor or labour market conditions. By plugging it into equation 1.41, under CARA higher wages lead to a higher effect of anticipated fertility on education.

$$a(T - \delta n)\tau P(I) > 1 \quad (1.42)$$

If we do not make the assumption of CARA, the effect is as ambiguous.

1.9.3.5 Thresholds $\bar{\alpha}$ and α

α is defined as, at $Z = 1$, equation 1.6 being the same for $n^* = 1$ and $n^* = 2$:

$$\begin{aligned} u'(c^2(n = 1)) \left[\frac{\partial p(I, n, Z = 1)}{\partial n} \Big|_{n=1} + \frac{\partial f(n)}{\partial n} \Big|_{n=1} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=1} \right] - \alpha \frac{\partial h(n)}{\partial n} \Big|_{n=1} = \\ u'(c^2(n = 2)) \left[\frac{\partial p(I, n, Z = 1)}{\partial n} \Big|_{n=2} + \frac{\partial f(n)}{\partial n} \Big|_{n=2} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=2} \right] - \alpha \frac{\partial h(n)}{\partial n} \Big|_{n=2} \end{aligned} \quad (1.43)$$

If $Z = 1$, the fines are 0 for the first and the second child ($n = 1$ and $n = 2$). Using this information and rearranging, gives:

$$u'(c^2(n=2)) \left[\frac{\partial f(n)}{\partial n} \Big|_{n=2} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=2} \right] - u'(c^2(n=1)) \left[\frac{\partial f(n)}{\partial n} \Big|_{n=1} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=1} \right] - \alpha \left[\frac{\partial h(n)}{\partial n} \Big|_{n=2} - \frac{\partial h(n)}{\partial n} \Big|_{n=1} \right]^{-1} \quad (1.44)$$

α does not depend on the fines.

$\bar{\alpha}$ is defined as, at $Z = 0$, equation 1.6 being the same for $n^* = 1$ and $n^* = 2$:

$$u'(c^2(n=1)) \left[\frac{\partial p(I, n, Z=0)}{\partial n} \Big|_{n=1} + \frac{\partial f(n)}{\partial n} \Big|_{n=1} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=1} \right] - \bar{\alpha} \frac{\partial h(n)}{\partial n} \Big|_{n=1} = u'(c^2(n=2)) \left[\frac{\partial p(I, n, Z=0)}{\partial n} \Big|_{n=2} + \frac{\partial f(n)}{\partial n} \Big|_{n=2} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=2} \right] - \bar{\alpha} \frac{\partial h(n)}{\partial n} \Big|_{n=2} \quad (1.45)$$

While $\frac{\partial p(I, n, Z=0)}{\partial n} \Big|_{n=1} = 0$, $\frac{\partial p(I, n, Z=0)}{\partial n} \Big|_{n=2} = \rho$. Therefore, we can reformulate:

$$u'(c^2(n=2)) \left[\rho + \frac{\partial f(n)}{\partial n} \Big|_{n=2} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=2} \right] - u'(c^2(n=1)) \left[\frac{\partial f(n)}{\partial n} \Big|_{n=1} - \frac{\partial y(I, J, n)}{\partial n} \Big|_{n=1} \right] - \bar{\alpha} \left[\frac{\partial h(n)}{\partial n} \Big|_{n=2} - \frac{\partial h(n)}{\partial n} \Big|_{n=1} \right]^{-1} \quad (1.46)$$

1.9.3.6 Gender Difference and Marriage Markets

The first order condition for equation 1.15:

$$\frac{\partial \omega(I^*, \tau)}{\partial I^*} \left(u'(c_m^2) \left[\frac{\partial y(I^*)}{\partial I^*} + \frac{\partial y(J(I^*), n)}{\partial J(I^*)} \sigma - \frac{\partial p(I, n, Z)}{\partial I^*} \right] + \alpha h(n) \right) + u'(c_s^2) \frac{\partial \tilde{y}(I^*)}{\partial I^*} \left(1 - \frac{\partial \omega(I^*, \tau)}{\partial I^*} \right) = \frac{\eta}{\delta} u'(c^1) \quad (1.47)$$

1.9.4 Further empirical results

1.9.4.1 Effect of fines on education

Denote P the official provincial fines. Once the full set of controls X is included, the effect of P is captured in the province times year fixed effects. However, we can look at the differential effect of fines for those that are eligible. Eligibility for a second child permit means an anticipated cost reduction in the monetary cost of the second child, and this cost reduction is a function of the fines. The cost reduction is essentially an interaction between eligibility and the monetary fine levels:

$$I_i = \delta \tilde{Z}'_i P_i + X_i + \tilde{\mu} \quad (1.48)$$

Monetary fines can also be used to test predictions derived from the theoretical model: First, monetary fines should change the distribution of *always-2* in the population. For this, I first check if eligibility has a higher effect on fertility outcome when fines are high. I also verify if monetary fines are important in making predictions in the prediction status *L*. Second, monetary fines do not have an effect for those that are either *increasers* and *always-1*. Third, monetary fines potentially have an effect on *always-2*, though it is unclear in which direction.

In order to investigate the effect of fines on education, I use the fine data set from [Ebenstein \(2010\)](#) which is based on [Scharping \(2013\)](#) and includes the official monetary fines (as well as bonuses) that are set on the province level. Unfortunately, the fines are only available until 2000, so that only 2/3 of the main sample can be used.²⁴ The fines are set as a multiple of the previous year's household annual income and ranged between 0.15 and 5 between 1990 and 2000. These official monetary fines measure the true costs of an out-of-quota child imperfectly: We do not know if these fines were actually implemented, and there were other penalties such as losing one's job that are important. Furthermore, there were several costly avoidance strategies whose employment probably correlated with the official fines.

The results are displayed in table 1.15. Second child permits can be interpreted as a cost reduction which equals eligibility status times the fine amounts. This definition of cost reduction is included in column 1 and 3 and the coefficients are interpreted as followed: a cost reduction of one time the annual income increases the likelihood of having a second child by 9.4 percentage points but does not have a significantly positive effect on education. Differentiating between those predicted to be *always-2* and those not (column 3), we find that those that are not predicted to be *always-2* react positively to a cost reduction while those that are predicted to be *always-2* do not. Here, the prediction process for *always-2* includes the fine levels as a possible predictor (see figure 1.8 on the bottom).

However, the amount of monetary fines does not add explanatory power over just using an eligibility dummy (see column 2 and columns 5 to 8). Being eligible in a province with higher fines does not have a higher effect on the likelihood of having a second child than being in a province with lower fines. Fines are also not important in predicting *always-2*. This is probably due to these official fines measuring the real cost of an out-of-quota child very noisily.

²⁴The previous results hold for this subsample.

Sample	Having 2nd child		Dependent variable: Enrollment in Senior High School			
	1977-1987	1977-1987	1990-2000	1990-2000	1990-2000	1990-2000
	(1)	(2)	(3)	(4)	(5)	(6)
Cost reduction	0.0941*** (0.0249)		0.0182 (0.0129)	0.0276* (0.0138)		
Eligibility at age 27/28		0.148*** (0.0433)				
Eligibility at age 27/28 X Monetary fine (std)		0.0114 (0.0110)				
Predicted <i>Always-2</i> X cost reduction				-0.0249** (0.0108)		
(Anticip.) eligibility at age 16					0.0929** (0.0431)	0.129*** (0.0397)
(Anticip.) eligibility at age 16 X Monetary fine (std)					-0.000107 (0.0239)	-0.00487 (0.0269)
Predicted <i>Always-2</i> X (Anticip.) eligibility at age 16						-0.0712 (0.0429)
Predicted <i>Always-2</i> X Monetary fine (std)						-0.0360 (0.0464)
Predicted <i>Always-2</i> X (Anticip.) eligibility at age 16 X Monetary fine (std)						0.0227 (0.0549)
Predicted <i>Always-2</i>				-0.0703** (0.0275)		-0.0706* (0.0379)
Observations	7607	7607	5405	5187	5405	5187
R^2	0.420	0.421	0.445	0.455	0.446	0.456

Note: All regressions are OLS Difference-in-Difference specification with province level times fixed effects, control for eligibility characteristics times province fixed effects and gender. Anticipated eligibility at age 16 is calculated on the basis of the province, eligibility characteristics and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. Column 1 and 2 control for sex of the first child, includes only those that are married and had their first child before 2003. Clustered standard errors on province X urban level in parenthesis. Significance levels: * 0.10; ** 0.05; *** 0.01 Data source: CFPS 2010.

Table 1.15: Monetary fines and eligibility.

1.9.4.2 Effect of second child permits on other variables

Anticipated fertility can also have an effect on other choices. For example, those that anticipate to have more children could get married earlier and have their first child earlier. However, in order to apply for a second child permit, couples often have to respect certain criteria, one being ‘late birth’ of the first child. Late birth is defined as the woman being at least 25 years old when having the first child. For those being born between 1970 and 1980, average age at first birth lies at 24.75 overall, 24 for women and 25.56 for men. Thus, this criterion might actually push those that plan to apply for a second child permit to wait until age 25 for having the first child.

Since the decisions are taken at a later stage in life, I use the anticipated eligibility status at age 22 for the timing of marriage, and the anticipated eligibility status at age 25 for the timing of first birth. The sample is adjusted accordingly. I also control for the years of education to correct for any effect that fertility expectations could have through education. The result using the triple differences specification are displayed in table 1.16. The coefficient of anticipated eligibility status at age 22 on the age of marriage is significantly negative once it is interacted with the sex of the individual (column 2). When female specific year and province fixed effects and an interaction between the female indicator and education are included (column 3), eligibility significantly decreases the age at marriage only for men.

We find the same results for the age at the birth of the first child: When female specific year, province, and educational effects are included, eligibility at age 25 significantly decreases the age at first birth only for men (column 6). A potential explanation is that eligible women are concerned with the official guidelines about late marriage and late birth. It is also possible that for women the legal limit for marriage is generally binding. At the same time, there are no thresholds for men and the legal marriage age might not be binding for most men, so that they can adjust the age of marriage and first birth to their anticipated fertility.

Sample (years they turned 16)	Dependent variable:					
	Age at marriage			Age at first birth		
	1981-1996 (1)	1981-1996 (2)	1981-1996 (3)	1984-1999 (4)	1984-1999 (5)	1984-1999 (6)
Nb children allowed at age 22	-0.618 (0.397)	-0.732* (0.419)	-0.949** (0.422)			
Nb children allowed at age 25				-0.680* (0.358)	-0.846** (0.376)	-0.967** (0.380)
Female X Nb children allowed at age 22		0.212 (0.276)	0.635*** (0.236)			
Female X Nb children allowed at age 25					0.314 (0.196)	0.602** (0.243)
Female dummy; Education	Yes	Yes	Yes	Yes	Yes	Yes
Eligibility controls	Yes	Yes	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes	Yes	Yes
Eligibility controls x Province FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE x Province FE	Yes	Yes	Yes	Yes	Yes	Yes
Eligibility controls x Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Female x Year FE,Province FE & Education	No	No	Yes	No	No	Yes
Observations	8703	8703	8703	9725	9725	9725
R^2	0.237	0.237	0.247	0.209	0.209	0.218

Note: Sample includes individuals that turned 22 between 1990 and 2000 for columns 1 to 3 and individuals that turned 25 between 1990 and 2000 for columns 4 to 6. Eligibility at age 22 and 25 is calculated on the basis of the province, eligibility characteristics and time. Eligibility controls include dummies for being an only child, having a rural household status and ethnic minority status. OLS regression with robust standard errors in parenthesis. Data source: China Family Panel Survey 2010.

Table 1.16: Effect of eligibility on age at marriage and age at first birth.

Chapter 2

Spousal Preferences and Parental Search: Data from China¹

EVA RAIBER, WEIWEI REN, JEANNE BOVET, PAUL SEABRIGHT,
CHARLOTTE WANG

¹This chapter is the result of a collaboration between the Institute for Advanced Studies in Toulouse - Toulouse School of Economics and the Yunnan Normal University. We are grateful for our student enumerators and all the support received from the staff at the Yunnan Normal University. We are thankful to Stéphane Straub, Emmanuelle Auriol, Abigail Adams, Matthias Doepke and discussants and seminar participants at the Development Workshop at the Toulouse School of Economics, the IAST Lunch Seminar and the 2017 IRMBAM conference sponsored by the IPAG Business School in Nice for their useful comments. The views expressed in this paper are solely those of the authors. Funding from the ANR - Labex IAST is gratefully acknowledged. Ren thanks the National Natural Science Foundation of China (71540032).

ABSTRACT

This chapter investigates marriage patterns and underlying preferences in China. We estimate mate preferences based on the evaluation of a series of randomly created profiles and connect our results to recent marriage patterns in the general population. Data is collected on parents (or other relatives) who search for a spouse on behalf of their unmarried adult child and on unmarried students. We confirm that male profiles with a high income and real estate ownership have a higher probability of being selected by parents than profiles with low income and no real estate, but not female profiles. We find that parents dislike profiles with less education than their son or daughter, but we do not find evidence for a dislike of female profiles with high education. However, some parents may have excessively high expectations about the educational level of their son-in-law. Parents mostly prefer male profiles with a similar age to their daughter and female profiles that are younger than their son. Yet, when their son is in his 20s, parents also accept women who are slightly older. If marriages formed according to parents' preferences, the most common case would be the husband being 2 or 3 years older. Yet, in the general population, the most common case for recent marriages is spouses having the same age. This would be the case if marriages formed according to students' preferences. Both parents' and students' preferences predict a high degree of assortativity on education, which we observe in recently married couples in the overall population.

2.1 Introduction

What do people look for in a marriage partner? Partner choice has important consequences for most life dimensions such as socioeconomic status, fertility, and overall well-being. It is therefore not surprising that in many contemporary societies (and most historical ones) parents are involved in the process of selecting a spouse for their child. While the economics literature on partner preferences has gained momentum (e.g. [Hitsch et al. \(2010\)](#), [Fisman et al. \(2006\)](#), [Xia et al. \(2014\)](#), [Ong and Wang \(2015\)](#), [Kurzban and Weeden \(2007\)](#), [Belot and Francesconi \(2006\)](#)), parental preferences have mostly been neglected in the research literature due to the focus on Western countries and the use of online dating data.² This chapter investigates parental preferences for spousal characteristics in China and compares them to the preferences of young adults. It then connects the results to recent marriage patterns.

Understanding parental and individual spousal preferences is not only interesting in itself but helps predict changes in the marriage structure due to, for example, imbalances in the sex ratio or to changes in the distribution of educational levels. Yet, we first need to understand if parental or individual preferences are driving marriage choices. Are parental preferences or individual preferences more reflected in the marriage outcome distribution? Furthermore, spousal preferences drive premarital investment choices. If educated women prefer educated men and parents prefer an educated daughter-in-law, it can encourage educational investment in sons. It is often policy-relevant to have estimates for these marriage market returns to education, real estate or age at marriage. For this, it is also crucial to understand if parents have a realistic assessment of their child's future marriage outcome. If parents overestimate marriage market returns, it can lead to prolonged search efforts and eventually lower marriage rates.

In China, the marriage prospects of women in their late 20s with a university degree are often described as uncertain in the media.³ Unmarried women over 27 were even labelled as “leftover women” by the government.⁴ Faced with a skewed sex ratio, lower marriage rates of educated women seems to be a policy concern for the Chinese government. Yet, [Brandt et al. \(2018\)](#) find that though university-educated women marry later, completed marriage rates are comparable to those of women without a university degree. Furthermore, it makes a difference whether lower rates of marriage among university-educated women are due to men (and their families) rejecting women who are more educated than them, or due to women (and their families) rejecting men who are less educated than them.

²Notable exceptions are [Banerjee et al. \(2013\)](#) and [Adams and Andrew \(2019\)](#) who focus on India.

³ see [Fincher, 2012, New York Times <https://www.nytimes.com/2012/10/12/opinion/global/chinas-leftover-women.html>](#)

⁴ see [Simpson, 2013, Mail online <https://www.dailymail.co.uk/news/article-2282249/The-leftover-women-China-defines-official-age-females-left-shelf-27.html>](#) and [Magistad 2013, BBC <https://www.bbc.com/news/magazine-21320560>](#)

Using the China Family Panel Study (CFPS) from 2014, a representative household survey, we find that only around 1% of those aged between 40 and 65 have never married. We also observe that recent marriages between 2012 and 2014 are highly assortative on age and education. Assortativity is characterized as spouses having the same educational level or similar age. Such assortative matches could be due to endogamic preferences, meaning preferences for those of the same type. They could also reflect increasing preferences - preferences for those with higher levels of education. If educational levels are similar between men and women, this leads to assortative matches.⁵ Finally, they can also arise if there are no specific preferences but individuals with similar characteristics have a higher probability of meeting. To be able to differentiate these options, we collected elicited preference data.

Parental influence over spouse choice continues to be important in China (Riley (1994), Pimentel (2000)), a country that had a tradition of arranged marriage before the Communist party came to power. In spring 2016, we interviewed men and women who were actively searching for a spouse for their adult child or relative at a public park in Kunming, in Southern China. The phenomenon of parents and other relatives searching for potential suitors in public parks at so-called “marriage markets” is nowadays common in big Chinese cities. Respondents were asked to evaluate a series of randomly created hypothetical profiles that we use to estimate spousal preferences. As a comparison group, we interviewed unmarried students at a local university.

We find that parents dislike profiles with an educational level that is lower than their daughter’s or their university- or college-educated son’s. Equivalently, female and male students dislike less educated profiles. Only parents searching on behalf of a university- or college-educated female subjects and female students have a clear preference for profiles with higher education. We find no evidence for a dislike of women who are more educated than the male eligible subject. High income and real estate ownership are important only for parents who search on behalf of their daughter, but not for those who search on behalf of their son. Parents searching on behalf of their daughter and female students prefer profiles with a similar age. Parents searching on behalf of men prefer younger women.

Our data set is of interest for several reasons. In contrast to data from online dating websites, we can credibly claim that respondents are looking for a spouse and not just for a short-term relationship. Furthermore, we are sure to have respondents who have thought about what they are looking for and are experienced in selecting potential candidates. Also, because choice data are taken from hypothetical profiles which are randomly created, we do not have the issue of a selected pool to choose from and we can ensure that the characteristics are not directly correlated. Since

⁵The highest type of man would then match with the highest type of woman, the second highest type with the second highest type and so forth. This holds under the assumption of non-transferable utility, which we make explicitly later in the chapter. It also holds under transferable utility under some assumptions on the combined marriage surplus (supermodularity).

we focus on economic variables and do not include physical attractiveness, results are less noisy and can be easily compared to outcome data from other data sets. The trade-off between income and physical attractiveness is investigated in a companion paper with a separate profile task (Bovet et al. (2018)).⁶

Popular belief is that people using this search platform are those who cannot find a spouse by conventional means. Comparing the search platform to representative household data from the same city, we find that the population represented at the platform is mostly local, older, more educated and more likely to be employed than the average unmatched population. This search platform seems to attract a specific subset of the unmatched population and it is not necessarily those with low quality. We also find that parental respondents rank altruistic motives, such as wanting their child to have a companion, as more important than selfish motives such as concerns for retirement and grandchildren.

Additionally to the profile choice task, we asked parents for the upper and lower limits for age and education that they find acceptable. The upper and lower limits give us “acceptable suitors”. We find that a match might not be accepted by the respondents because one party has too little education but not because one party is too educated. We investigate if recent marriages from the CFPS 2014 would be acceptable to our respondents. For age, we find that most marriages, in particular those between men and women in their early and mid-20s, are acceptable to parents. Yet, while parents with daughters in their late 20s and early 30s accept husbands with similar age, those with sons in their early 30s accept only younger women. For education, we find that 91% of wives have acceptable educational levels based on parents’ responses. Yet, only around 60% of husbands are within parents’ acceptable bounds. The partially unrealistic lower bounds on the educational level of a son-in-law could be an artefact of our selected sample. However, they could also be due to the (sometimes erroneous) expectation that their daughter’s spouse will have at least as much education as she has. With the educational achievements of women outperforming those of men in China in recent years, this could lead to lower marriage rates among educated women. However, there could also be an adjustment of the threshold over time as parents update their beliefs. Furthermore, we cannot rule out that respondent give us “strategic” lower bounds. Instead of stating what they would accept if there was no other option, these bounds could indicate the levels to which they aspire.

Based on parents’ and students’ preferences, we simulate marriages using a standard Gale-Shapley algorithm (Gale and Shapley (1962)). For the supply of spouses, we use individuals who recently married from the CFPS 2014 and compare the simulated outcomes with real outcomes. The baseline specification based on students’ and parents’ preferences predicts a higher degree of

⁶The profile task used in Bovet et al. (2018) took place after the hypothetical choice task used in this chapter. The results of Bovet et al. (2018) are in line with the results here: Respondent value income only when they are searching on behalf of a female subject. When they are searching on behalf of a male subject, income is insignificant.

assortativity on education than we observe in the real distribution. Frictions in the marriage sorting process, but also a higher heterogeneity in preferences in the general population compared to our sample, can explain this difference. The simulations based on parents' preferences predict the most common case to be couples where the husband is 2 or 3 years older. In the real distribution, the most common case is the spouses having the same age. This is the result when the simulation is based on students' preferences. The difference can be explained by changes in age preferences when the eligible subject, in particular a male eligible subject, gets older. It could, however, also reflect a difference in age preferences between generations. Matches between spouses of the same age seem to be preferred by students and accepted by parents. This hints at interesting dynamics of marriage decision-making and demands further research.

2.1.1 Links with the literature

This chapter adds to the growing strand in the economics literature that studies marriage patterns theoretically and empirically. The economic literature on marriage is based on the seminal work on marriage markets by [Becker \(1973\)](#). He models a market with two sides (men and women) and assumes transferable utility. Agents form matches and bargain about the distribution of the surplus that is generated by the match. The functional form of this marriage surplus is of particular interest. It reflects preferences of both sides. Following the seminal paper of [Choo and Siow \(2006\)](#), several papers have estimated the marriage surplus. It can be identified under relatively soft assumptions. However, without assumptions on the structural form, the bargaining and searching process, one cannot identify men and women's preferences separately. Notable papers in this field are [Choo and Seitz \(2013\)](#), [Wong \(2003\)](#), [Galichon and Salanié \(2015\)](#), [Logan et al. \(2008\)](#), and [Chiappori et al. \(2015\)](#). A summary of these methods can be found in [Chiappori and Salanié \(2016\)](#).

Another strand of this literature assumes non-transferable utility. Under non-transferable utility, men and women cannot bargain about the marriage surplus but each gets a determined fixed utility from the match. Most papers that assume non-transferable utility use additional data about preferences.⁷ This chapter follows the approach to identify preferences directly from additional data. Its empirical strategy is most closely related to [Hitsch et al. \(2010\)](#) who use data from an online dating website where they observe which profiles agents browse. They estimate which characteristics make a profile more likely to be clicked on. In this, they follow the searching framework of [Adachi \(2003\)](#). Similarly, [Banerjee et al. \(2013\)](#) use rankings of responses to marriage advertisements in an Indian newspaper. Other papers use data from speed dating ([Fisman et al. \(2006\)](#), [Fisman et al. \(2008\)](#)) or online dating ([Xia et al. \(2014\)](#), [Ong and Wang \(2015\)](#), [Kurzban and Weeden \(2007\)](#), [Belot and Francesconi \(2006\)](#)). [Adams and Andrew \(2019\)](#) use hypothetical marriage scenarios to elicit parents' spousal preferences for education and marriage age and beliefs about marriage prospects.

⁷A notable exception is [Del Boca and Flinn \(2014\)](#) who combine data on marriage outcomes and household production.

This chapter adds to the literature on the evolution of marriage patterns in China. [Mu and Xie \(2014\)](#) describe matching patterns on age over time in China. They find that homogamy in age, as measured by the average age difference between the spouses, increased until the 1990s, but since then has decreased slightly. We find similar results (see section 2.2.1). At the same time, the correlation between the age levels has increased continuously. [Hu \(2016\)](#) find that the household registration status (hukou) plays an important part in determining marriages. The household registration status is inherited from the parents and is an important determinant of socioeconomic status. [Du et al. \(2015\)](#) predict that due to the availability of more single men than women, women are more likely to marry ‘upwards’, in terms of individual income and hukou status.

Finally, this chapter touches on the issue of parents influencing the marriage outcomes of their children. [Huang et al. \(2016a\)](#) find that individuals that rely on their parents for finding a spouse end up with a different type of partner than if they found the spouse themselves. [Riley \(1994\)](#) uses data from a survey collected in 1986/1987. They find that the number of arranged marriages has decreased but point out that parents continue to have influence over marriage decisions.


2.2 Context

2.2.1 Marriage in China

The institution of marriage and the search for a marriage partner have undergone several changes in the last century in China. For centuries, arranged marriage had been dominant. Parents chose the spouse for their child, often with the help of a professional matchmaker ([Xia and Zhou \(2003\)](#), [Huang et al. \(2016a\)](#)). After the Chinese Communist Party came to power in 1949, the government passed a law making arranged marriages illegal (the Marriage Law, adopted in 1950). It was in line with its effort to get rid of traditional Chinese classes ([Engel \(1984\)](#)). Moreover, the government helped to abolish the traditional marriage system by encouraging women to join the labor force ([Pimentel \(2000\)](#), [Xia and Zhou \(2003\)](#)). However, in rural areas, arranged marriage continued to be important and parents still influenced marriage outcomes ([Xia and Zhou \(2003\)](#)).

The economic reforms of the late 1970s drastically changed the life of the Chinese people, as China became increasingly open to the rest of the world ([Higgins et al. \(2002\)](#), [Chang et al. \(2011\)](#)). The economic reforms led to an increase in economic opportunities, inequality and mobility. Marriage again became a way to increase social status and improve the financial situation ([Fan and Huang \(1998\)](#), [Han et al. \(2015\)](#)). This process was accompanied with changes in marital preferences. [Chang et al. \(2011\)](#) describes a decrease in the importance of virginity and an increase in the importance of financial prospects. [Higgins and Sun \(2007\)](#) and [Higgins et al. \(2002\)](#) find that though some attitudes are changing, attitudes toward sexual behaviors are still less open than in the West. They also note that students from cities have more liberal attitudes than those from

rural backgrounds.




Note: Answer to the question how the respondent met their spouse. Data source: China Family Panel Study 2010, marriages between 1945 and 2010. Omits categories with low incident rates (“other”) such as meeting online.

Figure 2.1: How individuals met their spouse.

Figure 2.1 shows the predominant ways in which individuals met their first spouse. The data used are from the China Family Panel Study (CFPS) 2010. We see a sharp drop in marriages that were arranged by their parents after the late 1940s and early 1950s. The share of those stating that they were introduced to their partner by a friend or relative increases at the same time, and this continues to be the predominant way. Unfortunately, we do not know who exactly introduced the spouses to each other, whether it was a common friend, a parent or another relative. We can also see that there is an increase in the proportion of couples that met at school/university and that met at work.

Chinese parents continue to influence their children’s marriage decision (Pimentel (2000)). As a sign for this, a new phenomenon appeared in several Chinese cities around 10 years ago: the so-called marriage markets. These platforms were created to help parents find a potential marital partner for their adult children. We will discuss these platforms in section 2.3.



Note: Educational level of husband minus educational level of wife. Educational levels: No/little education, primary school, secondary school, high school (normal and professional), tertiary education (university or college), graduate degree. Includes those that married between 2012 and 2014. Source: China Family Panel Study 2014

Figure 2.2: Educational difference between husband and wife


2.2.2 Marriage patterns in contemporary China

What do marriage patterns look like in China today? We use the China Family Panel Study (CFPS) 2014⁸, a nationally representative survey, which contains information on the age and educational level of married spouses. Restricting the sample to respondents between 20 and 65 years old, we find very high marriage rates. 82.7% of the respondents between 20 and 65 years are married, only 12.6% are single, 1.8% are divorced and 2.5% are widowed. For those between 40 and 65, only 1% has never been married (and 1.6% is divorced). We focus on those that got married in the last two years and are between 20 and 65 years old (910 observations). The average marriage age in this span is 26.8 for men and 25.2 for women.

As we can see in figure 2.2, marriages are assortative on the educational level. For nearly all educational levels, it is most common to be married to someone who has the same educational level (see also table 2.6 in the appendix). The only exception is for individuals who have a post-graduate degree, which is still relatively rare. The second most likely case is that the spouse has an educational level that is one below or above. The Pearson Correlation Coefficient between the educational level of husband and wife is highly significant at 0.65.

Furthermore, we investigate how spouses match along the dimension of age. Figure 2.3 displays the difference between the age of husband and wife. The average lies at the husband being 1.65 years older. However, the mode is at the wife and the husband having the same age. In most couples, the husband is up to three years older. In around 13% of the cases, the wife is one or two years older. There are some outlier couples in which the wife is more than five years older or the husband more than 10 years older. The Pearson's Correlation Coefficient is strongly significant at 0.85.

⁸The China Family Panel Study is a large-scale, nationally representative panel survey project conducted by the Institute of Social Science Survey at Peking University (information be found at www.iss.edu.cn/cfps/).

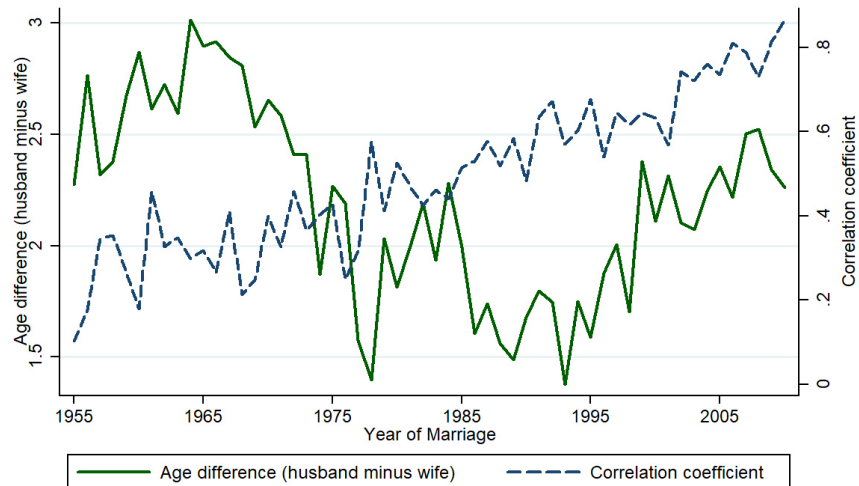


Note: Age of husband minus age of wife. Includes those that married between 2012 and 2014.
 Source: China Family Panel Study 2014

Figure 2.3: Age difference between husband and wife

We use the China Family Panel Study wave of 2010 to investigate how these age patterns have changed over time. We can see in Figure 2.4 that the age difference between a husband and a wife is higher in marriages that were formed between 1955 and 1970. Marriages that were formed between 1985 and 2000 have an age difference below 2 years and since then there was even an increase in the age differences. However, as the blue line indicates, the correlation coefficient between the husband's and the wife's age has been constantly increasing.

Two individuals get married if both of them (or their families) decide that the marriage is beneficial for them: marriage is an instance of a standard 1-to-1 matching procedure. Analysis of marriage data becomes difficult because both sides supply and demand spouses at the same time. Also, two (usually unmarried) individuals have to meet to consider marriage. The observed marriage patterns could be simply due to the chances of meeting. Educated men could be more likely to meet educated women. Men and women could be more likely to meet others of the same age. If they are older or younger, they might be already married or too young to get married. Yet, the observed marriage patterns can also be due to preferences. One might prefer a spouse with the same age or level of education (endogamic or homogamic preferences) or with a high educational level independently of one's own education. Without strong assumptions on search costs and the marriage surplus - the utility created by the marriage minus the utilities of being single - one cannot infer preferences from outcome data. This is why we have chosen to research partner preferences directly.



Note: Pearson’s correlation coefficient measuring the correlation between the age of husband and the age of wife. Includes those that married between 1955 and 2010. Source: CFPS 2010

Figure 2.4: Age difference and correlation between husband and wife’s ages over time

2.3 Parental Search at Green Lake Park

Every Saturday, one corner of the Green Lake Park in Kunming, the capital and largest city of the province of Yunnan in South China, hosts a “marriage market”.⁹ On this dedicated area of the park - accessible to the public - individuals search for a spouse for the “eligible subject”: either for themselves or for someone else.

This marriage search platform was initiated by parents who used their weekends to chat to other parents with unmarried children. Over time, it developed into an established event. Parents and other participants talk to each other or post papers with basic information of the eligible subject on the wall of the park, with their own contacts. They may check the information of others on the wall or address one of the marriage agencies present at the park. These marriage platforms are a new phenomenon but already widely spread. The local informal organizer indicated 2005 as the start of parents coming to the park with an explicit search motive. She also complained about the commercialization of the platform by the marriage agencies. The most famous marriage search platform in a public park is at the People’s Park in Shanghai is believed to have started in 2004.¹⁰

We ran a survey in this area of the park every Saturday from late spring to early summer of 2016¹¹ named the Questionnaire for Search Activities for a Marital Partner in Yunnan (QSAMPY). Our survey suggests that approximately 75% of all individuals who come to the park are looking

⁹With an estimated population of nearly 4 million (Cox (2018)), Kunming is a middle size city in China.

¹⁰According to an article in Al Jazeera: <https://www.aljazeera.com/indepth/inpictures/2013/04/201343113125739211.html>.

¹¹With some exceptions due to the exam period.

for a marriage partner on behalf of someone else. Around 23% look for a spouse for themselves. Those that are looking for themselves are older than those who are represented by their parents (48 years on average) and have mostly been married before. Those that are represented by their parents, an uncle or an aunt, are younger (on average 33 years old) and usually unmarried (94%). We are interested in preferences for the first marriage and of those who search on behalf of someone else. We thus exclude those searching for themselves. We call this the parents' sample, though it includes other relatives if not explicitly specified.

2.3.1 Who is searching?

Respondents who are looking on behalf of someone else are predominantly female (64%) with an average age of 61. The majority of respondents (86%) are retired. This suggests that they have low opportunity costs of going to a park on a Saturday. The vast majority (91%) are married, 5% are widowed and 2.5% are separated or divorced. Only 5% of respondents are at the park together with the eligible subject. Around half of the respondents are searching on behalf of their daughter, one third on behalf of their son, and the rest on behalf of their niece, their nephew or another relative. Summary statistics are displayed in table 2.7 in the appendix.

The park where the data were collected is in the centre of the city. Most of the respondents of the QSAMPY live and are registered in the city. The search platform does not attract many non-locals, for example from rural areas. Of parents, 82% live with the eligible subject. As expected for an urban Chinese sample, most parents only have one child (73%). 23% have two children and 4% have three or more children. In 75% of the cases when the respondent has more than one child, the other child is already married. In 68% of these cases, the other child already has at least one child.

The province of Yunnan is ethnically diverse. However, most respondents (94%) state that they are of Han ethnicity. The share of ethnic minorities in our sample is thus lower than in the city overall, where it was around 15% in 2007.¹²

Do respondents search on behalf of someone else because they themselves were introduced to their spouse by someone else? We find that 36% of QSAMPY respondents were introduced to their spouse by friends. 30% met at work or in the collective unit, 6% met at school or university and 13% were introduced by a relative (not their parent). The share of those being introduced by their parents only lies at 3%. Only 3 respondents (less than 1%) in our survey indicate that their marriage was arranged (by their parents, relatives or a marriage agency). These numbers are comparable with those of the general population displayed in figure 2.1.

¹²The original source for this information is not available anymore. However, it is quoted in several other websites, such as Wikipedia and https://www.gokunming.com/en/blog/item/397/kunming_residents_by_the_numbers.

2.3.2 How are they searching?

A large share of respondents comes to the park occasionally (35%), some regularly (18%), and an important share always (22%). 23% indicate that this is the first time they are at the park. Parents (and other relatives) invest an important share of their Saturday on the search. A quarter of respondents spends between 1 and 2 hours at the park. Only 4% spend less time, 21% spend 2-3 hours, 17% 3-4 hours, 17% 4-5 hours and 13% spend more than 5 hours. Respondents focus on finding someone at this platform: only 4% use another platform. Of those, most use online platforms.

Several marriage agencies are present at the park. They display big folders that are supposedly full of information on unmarried potential spouses. They sell the contact information at a piece rate (or several for a certain price) and ask a fee for being included in the database.¹³ However, we find that most parents (83%) have not paid anything to a marriage agency and only 2.7% of eligible subjects ever went to a meeting arranged by a marriage agency.

Most of the eligible subjects (72%) know that they are represented at the park. Of those, the majority (68%) approves of this procedure. 30% of the eligible subjects encourage their parents or other relatives to search on their behalf. Nearly half of the respondents (42%) indicate that the eligible subject is also searching actively for a spouse. In the case of parents, the spouse of the respondent is sometimes also involved in the search process (36%). In the case of another relative, the parents are often looking as well (67%). Finding a spouse for the eligible subject often seems to be a family affair where more than one person is involved.

2.3.3 Why are they searching?

Nearly half of the respondents (42%) indicate that the eligible subject is also searching actively for a spouse. Of those that state that the eligible subject is not searching, ‘not having time/being too busy’ is the most stated reason. It is followed by ‘not wanting to’ and ‘being too shy’.

We offered pre-selected reasons for the search to investigate the respondents’ search motivations. We asked if they thought this reason was important or not for them (on a Likert scale of 5). We find that parents rank altruistic motives as more important (see table 2.8 in the appendix). These include ‘wanting their child to have someone who takes care of him or her’ and ‘having their child to have someone he/she feels affection’ for. More selfish reasons are less often ranked as important or very important independent of the sex of the eligible subject. These included ‘having a son or daughter-in-law take care of the respondent when the respondent is older’ or ‘having grandchildren’.

¹³We talked to two marriage agencies and visited their office. For one agency, to be included in the database, one has to come to the office with documents to verify the information provided and to take a picture. The other agency just had basic information and no pictures. We were not able to verify the reliability of their statements.

These survey results suggest that parents see themselves as agents for their children. They want to help their children by searching on behalf of them. This does not exclude that parents put direct or indirect pressure on their children to find a spouse.

2.3.4 On whose behalf are they searching?

In the QSAMPY, we collected detailed information on the “eligible subject”. It includes age, sex, marital status, education, income, ethnicity, religion, physical characteristics and habits. Table 2.9 in the appendix summarizes some descriptive statistics. The average age of eligible subjects is 33 years. There are more women being represented than men (58%). The sample is very educated: 98% have at least finished professional high school, 60% have completed a university degree. 21% of eligible subjects are represented by their father, 64% by their mother. The rest is represented by an uncle or an aunt, in few cases by another relative or a friend. Among the eligible subjects, 92% are employed and 5% are self-employed or entrepreneurs. income is 4800 Yuan; most individuals earn between 2000 and 8000 Yuan. 94% of eligible subjects are of Han ethnicity, 2.8% of Bai ethnicity, and 2.2% of another ethnicity.

We use the roster data from the Skills Towards Employability and Productivity (STEP) program to compare our sample to the general population. The STEP data in Yunnan, China, was collected in 2012 by the World Bank within urban Kunming. The sampling method follows the census paths to be representative. We only use the first step of the questionnaire, the household roster. It contains basic information about all the individuals living in the household. We use the data about the sex, age, marital status, education and income of those who are at least 20 years old.

In table 2.1, we compare the population from the QSAMPY and the STEP data. We find that never married eligible subjects at the park are significantly older than the average never married person (by around 6 years). Those that have no partner (including divorced and widowed) have a similar average age to those at the park. The sex ratio in the general population is balanced. Yet, the percentage of women who never married or are without a partner is only 44% in the STEP data. This may reflect a skewed sex ratio (though Yunnan has one of the most balanced sex ratios in China). Another more plausible reason is that women on average marry earlier. The comparison with the share of women at the park indicates that women are over-represented in the QSAMPY data.

The QSAMPY sample is much more educated than the general population. 98% of eligible subjects have finished at least professional high school and around 60% have a university degree. In the STEP data, only 1 person completed tertiary education and only 55% of unmarried individuals finished high school. Economic indicators are also better in the QSAMPY. 97% of the eligible

¹⁴STEP: Has been working at least 1 hour in last 7 days. QSAMPY: Employed or self-employed

	QSAMPY (2016)		STEP (2012)			Difference		
	Eligible Subjects		never married	no partner	all	never married	no partner	
	never married	no partner						(1)
Age (mean)	32.43	33.06	26.19	33.70	44.40	6.24***	-0.64	
Female	59%	58%	44%	48%	51%	15pp***	10pp***	
Completed min. (prof.) high school	98%	98%	55%	46%	25%	43pp***	52pp***	
University degree	62%	60%	0%	0%	0%	62pp***	60pp***	
Working ¹⁴	97%	97%	65%	57%	62%	32pp***	40pp***	
N	377	410	516	649	4970			

Column (1) and (3) include information on those who have never married, column (2) and (4) who currently do not have a partner (those unmarried/divorced/widowed and single), column (5) includes the whole population (including married and with partner). Asterisks indicate significant difference at 1% (***). Comparisons are: QSAMPY never married to STEP never married; QSAMPY no partner - STEP no partner. Two sample t-test. Data sources: QSAMPY 2016 and STEP 2012.

Table 2.1: Comparison between STEP and QSAMPY data

subjects are employed or self-employed. In the STEP sample, only 62% indicate having worked at least one hour in the last seven days. This percentage goes up only slightly if we only look at those that have never married.

There are obvious limitations to the comparability between the QSAMPY data collected in 2016 and the STEP data in 2012. The marriage age between 2012 and 2016 presumably increased, as did educational attainments. However, the difference in the average age is unlikely to just be driven by an increase in the marriage age. Also, the striking difference in educational levels is very unlikely to be driven by a general increase in education. We argue that those represented by their parents at Green Lake Park are older than the average unmarried population. They also have a higher social status measured in educational levels and employment rates. The Green Lake Park search platform seems to be a specialized platform attracting educated working individuals that are above the average marriage age.

2.3.5 Comparison: The student sample

As a comparison group, we collected data from 283 students at the Yunnan Normal University. Students were randomly approached in front of the student canteens. Students are commonly used for surveys about marriage preferences because they are generally unmarried but of marriageable age. The results obtained from this sample can easily be compared with other studies. Also, one might wonder if parents' preferences are actually decisive for the marriage choice. Do they differ from those of their adult children? Comparing students' preferences to parents' preferences shows

where they overlap and where not.

The student sample has obvious drawbacks. First, the students are not the children of the parents we interview but a distinct sample with different characteristics. Second, students are not necessarily looking for a long-term partner at the time of the interview. Thus, they might not have very clear preferences or confound them with preferences they have for a current short-term partner. This is a common drawback of stated preferences.

The student sample is female biased: 57% of the respondents are women. However, this reflects the university population where women are over-represented. The average age is 21. The student sample is more ethnically diverse. 82% of the students are Han Chinese, 5% Bai, 5% Yi and 8% from other ethnicities. Only 11% of the students are from the city of Kunming. The majority (65%) are from another place in Yunnan and 23% from another province. The student sample is, therefore, more heterogeneous in terms of ethnicity, origin and religion. At the same time, it is more homogeneous in terms of age and education.

2.4 Preference Estimation

2.4.1 Hypothetical profiles

During the interviews, respondents were shown four pairs of hypothetical profiles. The profiles had information about age, income, education, ethnicity and real estate ownership. We asked respondents to state if the profile represents a person that they would want their eligible subject to meet (“meet choice”). Afterwards, they were asked which profile of the two they preferred (“preference choice”). For the “meet choice”, we have 8 observations per respondent. For the “preference choice”, we have 4 observations per respondent. In this chapter, we focus on the “meet choice”, however, the preferences derived from the “preferred choice” are essentially the same (compare the results to table 2.10 in the appendix). The profiles were clearly described as hypothetical and the answers were not incentivized. The choice data, thus, do not provide revealed preferences. However, respondents are about to choose who they want their eligible subject to meet at the park. They have thought about which characteristics they prefer and are about to evaluate the same information. They do not have an incentive to deviate from the strategy they use for their actual choices. Therefore, stated preferences are presumably close to revealed preferences in this context.

The advantage of the hypothetical choice data is that characteristics are randomly created. In other studies, the choice set is restricted due to platform entry or first sign of interest (e.g. Banerjee et al. (2013) and Hitsch et al. (2010)). The only dimension the choice set was restricted is age. Before the respondents saw the profiles, they could choose the age category they were interested in. Respondents had to choose between the ranges: 20 to 39 year old, 30 to 49 years old, 40 to 59

years old and over 50 years old. The age brackets each include 20 years or more. On the one hand, we wanted to avoid repeatedly showing participants profiles that were far off the desired age range. On the other hand, we wanted to observe until which age difference they found profiles acceptable. Thus, the age brackets encompassed 20 years or more.

Figure 2.5 displays an example profile pair. The educational level was drawn from junior high school degree (junior middle school), senior high school degree, university (bachelor) degree and graduate (masters degree).¹⁵ High school degree and university degree had a higher likelihood of being drawn for the creation of the profile. This was designed to avoid participants' facing several unacceptable profiles. The indicated monthly income indicated between 2000 Yuan and 8000 Yuan. The ethnicity was either Han, Yi or Dai, with Han ethnicity having a higher likelihood to be on the profile. Real estate property indicated either "yes" or "no" with the same likelihood and was not further specified. The profiles mimic information that participants usually share with others at the park via written CV.

Profile A		Profile B	
Age:	35	Age:	30
Education:	Bachelor Degree	Education:	High School
Monthly Income:	2000	Monthly Income:	4000
Ethnicity:	Han	Ethnicity:	Bai
Real Estate:	Yes	Real Estate:	Yes

Figure 2.5: English example for the hypothetical profiles that were shown to respondents.

2.4.2 Estimation Framework

In theory, participants agree to a meeting if they expect the utility generated by the potential match to be higher than their "reservation utility" which is the utility of staying single and continuing the search. We define the utility a woman w gets from marrying a man m as $u_w(m)$ and the utility a man m gets from marrying a woman w as $u_m(w)$. We denote the reservation utility of women w as v_w and the reservation utility of man m as v_m . We can then define the probability of the respondents indicating that they would want their subject to meet the person described in the profile as the probability that the utility derived from the match is higher than the reservation utility:

$$Prob(Meet = 1) = Prob(u_w(m) - v_w \geq 0) \quad (2.1)$$

¹⁵Education is officially mandatory until the completion of junior high school (9 years of education). Afterwards, students can decide to continue schooling at different types of senior high school (regular or vocational) for three years which was not further specified in the profiles.

This method is derived from the model of Adachi (2003) and also used in Hitsch et al. (2010). We can parametrize the utility function of woman w $u_w(m)$ as a function of the man m 's observed characteristics, how they interact with woman w 's characteristics plus an unobserved term that captures unobserved characteristics including how well the two get along:

$$u_w(m) = f(m, w'm, \theta) + \epsilon_{w,m} \quad (2.2)$$

with θ being the parameters to estimate.

Assuming $\epsilon_{w,m}$ is i.i.d. with the standard logistic distribution, we can derive the binomial logistic regression:

$$Prob(Meet = 1) = Prob(f(m, w'm, \theta) - v_w \geq 0) = \frac{\exp(f(m, w'm, \theta) - v_w)}{1 + \exp(f(m, w'm, \theta) - v_w)} \quad (2.3)$$

This equation can now be estimated using a logit regression. We control for v_w/v_m by including individual fixed effects since we have repetitive choices. However, this implies that respondents that always stated either 'yes' or 'no' to all profiles, are discarded from the estimation.

2.4.3 Results: Parents

We run a logit regression with the dependent variable being the meeting indicator. The indicator equals 1 if the respondent would want the eligible subject to meet the person described in the profile. The results are displayed in table 2.2. We look at men and women separately. At first, the whole sample is included. We then split the sample between those where the eligible sample does and does not have a university degree (including a college degree which is more applied). The reference category for education is "Profile: Only Junior High School". We include individual fixed effects and cluster standard errors at the individual level. The coefficients display average marginal effects.

In table 2.3, the educational levels and income are replaced by the difference between the eligible subject's educational level/income and the profile's educational level/income. There are some missing values particularly for the eligible subject's income. Thus, this specification has fewer observations. There is too little variation in the ethnicity of the eligible subject (most of them are Han Chinese). We can unfortunately not distinguish if individuals prefer a spouse with Han ethnicity or with the same ethnicity.

Dependent variable: Indicator of wanting to meet								
	Female eligible subject				Male eligible subject			
	all (1)	w/o degree (2)	with degree (3)	all (4)	all (5)	w/o degree (6)	with degree (7)	all (8)
Profile: HS degree	0.117*** (0.0306)	0.141** (0.0717)	0.113*** (0.0319)	0.114** (0.0444)	0.129*** (0.0458)	0.107+ (0.0658)	0.208*** (0.0499)	0.115** (0.0460)
Profile: Undergraduate degree	0.304*** (0.0340)	0.240*** (0.0717)	0.334*** (0.0424)	0.301*** (0.0641)	0.161*** (0.0472)	0.0650 (0.0696)	0.314*** (0.0589)	0.147*** (0.0476)
Profile: Graduate degree	0.362*** (0.0436)	0.228*** (0.0784)	0.416*** (0.0502)	0.358*** (0.0688)	0.123** (0.0559)	-0.0614 (0.0768)	0.392*** (0.0713)	0.0983* (0.0549)
Profile: Han ethnicity	0.0507** (0.0227)	0.0985** (0.0451)	0.0265 (0.0256)	0.0490** (0.0223)	0.0342 (0.0313)	0.00970 (0.0381)	0.0547 (0.0466)	0.0268 (0.0311)
Profile: Log(income)	0.163*** (0.0287)	0.178*** (0.0533)	0.146*** (0.0345)	0.161*** (0.0304)	-0.0320 (0.0405)	-0.109** (0.0527)	0.0754 (0.0584)	-0.0366 (0.0406)
Age difference (+)	-0.0236*** (0.00416)	-0.0171** (0.00714)	-0.0259*** (0.00514)	-0.0202*** (0.00736)	-0.0481*** (0.00930)	-0.0528*** (0.0144)	-0.0448*** (0.0123)	-0.0499*** (0.0109)
Age difference (-)	-0.0454*** (0.00684)	-0.0249** (0.0106)	-0.0545*** (0.00970)	-0.0711*** (0.0121)	0.0102* (0.00525)	0.00609 (0.00731)	0.0103+ (0.00666)	0.0567*** (0.0143)
Profile: Owns real estate	0.0724*** (0.0260)	0.0842* (0.0452)	0.0695** (0.0309)	0.0742*** (0.0236)	0.0282 (0.0294)	0.0286 (0.0391)	0.0449 (0.0415)	0.0247 (0.0283)
Age difference (+) squared				-0.000437 (0.000474)				0.00127*** (0.000338)
Age difference (-) squared				0.00284*** (0.000771)				-0.00376*** (0.00102)
Observations	1726	502	1224	1726	1184	656	528	1184

Note: Logit regression including individual fixed effects, coefficients indicate average marginal effects, standard errors, clustered on the individuals level, are in parenthesis. Female eligible subject = respondent is choosing on behalf of a woman; male eligible subject = respondent is choosing on behalf of a man. Reference category for education: "Profile: Only Junior High School". "HS" is the abbreviation for high school. Age difference is calculated at the age of the profile minus the age of the eligible subject. Source: QSAMPY 2016

Table 2.2: Estimation using hypothetical choice data and the question which profile the respondent would like their eligible subject to meet.

Dependent variable: Indicator of wanting to meet

	Female eligible subject			Male eligible subject		
	all (1)	w/o degree (2)	with degree (3)	all (4)	w/o degree (5)	with degree (6)
Educational difference (+)	0.0437 (0.0305)	0.00530 (0.0429)	0.0937** (0.0453)	-0.0444 (0.0318)	-0.0473 (0.0375)	0.0434 (0.0597)
Educational difference (-)	-0.176*** (0.0193)	-0.154*** (0.0484)	-0.173*** (0.0213)	-0.119*** (0.0273)	-0.0420 (0.0476)	-0.149*** (0.0301)
Profile: Han ethnicity	0.0461** (0.0227)	0.0875* (0.0457)	0.0240 (0.0256)	0.0319 (0.0316)	0.0106 (0.0406)	0.0497 (0.0481)
Income difference (+)	0.130** (0.0521)	0.177** (0.0902)	0.0831 (0.0684)	-0.0595 (0.0727)	-0.114 (0.0928)	0.0715 (0.119)
Income difference (-)	-0.190*** (0.0465)	-0.171 ⁺ (0.107)	-0.192*** (0.0503)	-0.00854 (0.0609)	0.0944 (0.0916)	-0.0802 (0.0773)
Age difference (+)	-0.0233*** (0.00421)	-0.0171** (0.00707)	-0.0259*** (0.00520)	-0.0500*** (0.00954)	-0.0515*** (0.0148)	-0.0480*** (0.0131)
Age difference (-)	-0.0445*** (0.00709)	-0.0213** (0.0105)	-0.0547*** (0.00971)	0.00865* (0.00522)	0.00495 (0.00747)	0.0106 ⁺ (0.00695)
Profile: Owns real estate	0.0703*** (0.0263)	0.0865* (0.0451)	0.0670** (0.0309)	0.0326 (0.0299)	0.0222 (0.0408)	0.0508 (0.0420)
Observations	1694	486	1208	1152	640	512

Note: Logit regression including individual fixed effects, coefficients indicate average marginal effects, standard errors clustered on the individual level are in parenthesis. Female eligible subject = respondent is choosing on behalf of a woman; male eligible subject = respondent is choosing on behalf of a man. The differences are calculated at the characteristic of the profile minus the characteristic of the eligible subject. Source: QSAMPY 2016

Table 2.3: Estimation using hypothetical choice data and the question which profile the respondent would like their eligible subject to meet.

2.4.3.1 Results for Women

We first have a look at spousal preferences when the eligible subject is a woman. The respondents are thus looking for a son-in-law.

- **Education:** The likelihood of wanting a meeting increases significantly with the educational level of the profile. This holds for both female subjects with and without a university degree (table 2.2 column 1 to 3). On average, a profile with a graduate degree is 36 percentage points more likely to be chosen to meet than one with only junior high school. The marginal effect of education is higher when the female subject has a university degree herself. We find that respondents dislike male profiles with a lower educational level (see table 2.3). When

the female eligible subject has a university degree, respondents have a preference for male profiles with higher education.

- **Income and Real Estate:** Respondent prefer male profiles with higher income. A profile with double the income has a 16 percentage point higher likelihood to be selected. The coefficient is higher for women who do not have a university degree but not significantly so. In table 2.3 we see that for respondents with a university-educated female subject, it is important that the man earns as much as her, but not necessarily more than her.

Owning real estate also increases the likelihood of being selected by around 7 percentage points. This holds for female subjects with and without a university degree.

- **Age:** Respondents dislike a negative and a positive age difference between the eligible subject and the profile. They prefer someone of a similar age. We include squared terms of the age variables (table 2.2 column 7). Figure 2.6 shows the average predicted likelihoods. Respondents dislike a negative age difference more than a positive age difference. For a positive age difference, the predicted likelihood is slowly decreasing (only 2 percentage points per year).




Figure 2.6: Female subjects: The predicted likelihood of choosing the profile according to the age difference (age of profile minus age of eligible subject).

- **Ethnicity:** Only respondents choosing on behalf of a female subject without a university degree have a preference for Han ethnicity. They are around 10 percentage points more likely to select a profile that states “Han ethnicity” than one of the minority ethnicities.

2.4.3.2 Results for Men

We now look at spousal preferences when the eligible subject is a man. Respondents are thus looking for a daughter-in-law.

- **Education:** When the male subject has a university degree respondents prefer educated female profiles. For university-educated subjects, a profile with an undergraduate degree has

a 31 percentage point higher likelihood to be selected than one that has only a high school degree (table 2.3 column 6). This does not hold for subjects that do not have a university degree. Educational preferences for subjects without tertiary education are insignificant. We do not find evidence for a dislike of “too educated” women. The coefficients for higher and lower education that the subject has are negative, but far from being significant (table 2.3 column 6). For men with a university degree, respondents dislike profiles with a lower educational degree. They do not have significant preferences for a higher degree either.

- **Income and Real Estate:** Respondents choosing on behalf of subjects without a university degree have a dislike for high income. For those choosing on behalf of men with a university degree the coefficient is positive but not significant. Looking at the relative income, none on the coefficients are significant. It is still interesting that the signs are opposite. For men without a university degree, respondents might be less likely to select the profile when the woman earns more than the man. This might be the opposite for men with a university degree. There is no significant preference for the spouse owning real estate.
- **Age:** Respondents dislike a positive age difference where the woman is older than the eligible subject. Introducing squared terms (column 8), we find that respondents prefer younger women but this preferences decreases with the age difference. Looking at figure 2.7, we see the dislike of (much) older woman on the right panel and a preference for younger women on the left panel. The mathematically optimal age difference based on the coefficient lies at 15.08 years. Yet, we have less power at higher age differences and do not know if preferences are decreasing afterwards.




Figure 2.7: Male subjects: The predicted likelihood of choosing the profile according to the age difference (age of profile minus age of eligible subject).

- **Ethnicity:** The variable for ethnicity is not significant.

2.4.4 Results: Students

The results from the student sample are displayed in table 2.4. Since the students have not entered the labour market, we do not know how much income they will earn. Furthermore, we do not know their realized educational level. We therefore only run the estimation with income and education level, but not differences.

- **Age:** Female students have a preference for meeting someone with a similar age. They have a clear dislike of meeting a man who is younger than themselves. Male students have a dislike for meeting someone who is older. They seem to be indifferent between someone who has the same age and someone who is younger.
- **Education:** Both male and female students have a strong preference for someone who is educated. The likelihood of the profile being chosen for a meeting is the highest with a graduate degree when a woman is choosing and an undergraduate degree when a man is choosing.
- **Ethnicity:** Both male and female students have a preference for Han Chinese.
- **Income and Real estate:** Both male and female students have a preferences for higher income. Interestingly, the coefficient for male students is nearly as high as the coefficient for female students. Female students have a preference for someone who has real estate.

2.4.5 What do we learn from the preferences?

In our dataset, we find that respondents have preferences for male profiles with at least as much education as the female subject. For educated female subjects, respondents clearly have increasing preferences: the more education the better. We see the same patterns among female students. Respondents searching on behalf of men with a university degree also prefer educated profiles. The coefficient for a negative educational difference is high. This making education one of the key characteristics. We see the same among male students, who place high importance on female profiles having a university degree. Yet, when the man does not have a university degree, education is not important. Overall, we do not find evidence for a dislike of “over-educated” women.

Preferences for age are endogamic for respondents searching on behalf of a female subject and female students: they prefer male profiles with a similar age. Yet, they are more tolerant of the man being older than younger. Male students and respondents have a dislike against the female profile being older. There is weak evidence that they favour younger female profiles. Yet, the coefficients are only clearly significant in the non-linear form in the pooled parents sample.

Thus, even though there is no explicit dislike of educated women, there might be a weak preference against the combination of being educated and older. Consider a respondent choosing between

Dependent variable:
Indicator for selecting 'meet' option

	Female subject			Men subject		
	(1)	(2)	(3)	(4)	(5)	(6)
Profile: HS degree	0.177*** (0.0427)	0.178*** (0.0427)		0.121*** (0.0456)	0.118*** (0.0457)	
Profile: Undergraduate degree	0.371*** (0.0450)	0.371*** (0.0449)		0.376*** (0.0499)	0.374*** (0.0501)	
Profile: Graduate degree	0.429*** (0.0520)	0.433*** (0.0518)		0.297*** (0.0590)	0.303*** (0.0589)	
Profile: Han ethnicity	0.140*** (0.0271)	0.137*** (0.0269)	0.140*** (0.0271)	0.101*** (0.0321)	0.105*** (0.0319)	0.0926*** (0.0328)
Profile: Log(income)	0.215*** (0.0328)	0.215*** (0.0328)	0.217*** (0.0326)	0.154*** (0.0335)	0.155*** (0.0334)	0.148*** (0.0337)
Age difference (+)	-0.0263*** (0.00427)	-0.0103 (0.0122)	-0.0265*** (0.00427)	-0.0457*** (0.00709)	-0.0670*** (0.0150)	-0.0444*** (0.00724)
Age difference (-)	-0.0757*** (0.0127)	-0.0342 (0.0445)	-0.0764*** (0.0126)	-0.0115 (0.00882)	0.00687 (0.0259)	-0.0115 (0.00898)
Profile: Owns real estate	0.0665** (0.0305)	0.0655** (0.0306)	0.0625** (0.0303)	0.0374 (0.0393)	0.0379 (0.0396)	0.0337 (0.0392)
Age difference (+) squared		-0.000957 (0.000750)			0.00194* (0.00112)	
Age difference (-) squared		-0.00665 (0.00705)			-0.00380 (0.00324)	
Educational difference (+)			0.105*** (0.0301)			0.0130 (0.0379)
Educational difference (-)			-0.199*** (0.0263)			-0.205*** (0.0258)
Observations	1168	1168	1168	930	930	930

Note: Logit regression including individual fixed effects, coefficients indicate average marginal effects, standard errors are in parenthesis. Source: QSAMPY 2016

Table 2.4: Student Sample: Estimation using hypothetical choice data and the question which profile the respondent would like to meet.

a woman with a high school degree and a woman with a university degree with the same age as the male subject. Everything else equal, the average respondent would choose the woman with the university degree. This advantage can be offset if the less educated woman is ten years younger than the male subject. However, if the university-educated woman was a couple of years older than the male subject, the less educated women would only have to be two years younger than her to make up for the educational disadvantage. The trade-off is stronger for male subjects without tertiary education, as the respondents do not place any significant value on education.

Respondents and students generally prefer high-income male profiles with real estate (table 2.2). For female subjects without a university degree, preferences for income are increasing. For university-educated female subjects, respondents prefer men who earn at least as much income as the female subject. Earning more than her is not as important. We cannot verify this preference among female students as they have not yet entered the labour market. We find support for the idea that for men having real estate is important in the marriage market. However, the marginal effects on education are much higher. A university degree compared to junior high school increases the likelihood of being selected by parents by 30 percentage points and 37 percentage point by students. Senior high school compared to junior high school, increases the likelihood by 11 percentage points (18 percentage points for students). Having real estate increases the chances to be selected by only 7 percentage points.

On the other side of the market, respondents searching on behalf of men without a university degree dislike female profiles with high income (table 2.2). Combined with table 2.3, it seems that the dislike is only towards those who earn more than the male eligible subject. Yet, for men with a university degree and male students, we do not find this preference.

Overall, we find that educated women with high income have good chances among those that they themselves prefer: men with education. There is a weak preference among respondents searching on behalf of men without a university degree against women who earn more. Yet, those searching on behalf of such women have a stronger preference against such men. The same holds for education. Respondents searching on behalf of men without a university degree do not place any significant value on tertiary education. Yet, those searching on behalf of university-educated women and female students have a strong preference against these less educated men. For men with a university degree and male students - the “target group” of respondents searching on behalf of a female subject and of female students - respondents place value on education and show no signs of a dislike for the women being more educated or having more income.

2.5 Comparison with marriage patterns

The estimated preferences offer interesting insights into the likes and dislikes of parents and young adults. However, can they explain the marriage patterns in the general population? This question is crucial as our two samples are not representative of the general population. What claims can we make about marriage preferences in the general population? Furthermore, we do not know how much influence parents and individuals have over mate choice. Are parents' or students' preferences more reflected in the general marriage pattern?

To address these questions, we connect our preference and questionnaire data to the recent marriage patterns in the general population. We first look at the accepted levels of education and age. We then simulate marriages based on the estimated preferences and compare simulated marriages with real marriages.

2.5.1 Stated acceptance limits

Respondents at the park were asked which level of education and age they would accept for their daughter- or son-in-law. For these variables, they indicated a lower and an upper bound of what would be acceptable. They could also choose “no limit”. The lower and upper bounds then give us acceptance sets for men and women: all those that lie between the limits are “acceptable” to parents. When the acceptance sets overlap, matches are agreeable by both parties and thus feasible. The area where acceptance sets overlap is called the correspondence set. Students also answered these question - however, they are of less interest because students are homogeneous in age and education. We focus on how the acceptance limits for age and education and investigate how they vary with the age and education of the eligible subject.

Asking respondents about their acceptable limits is similar to asking individuals about their willingness-to-accept for an object without incentivizing the answer. Yet, in our case, respondents are about to select potential suitors in real life and this question asks them about the decision rule that they are just now employing. However, the answers should not be interpreted as the limits if there is clearly no other option, but as the cut-off that they are currently employing which may change over time. Therefore, they reflect also what they aspire their son- or daughter-in-law to be like.

The acceptance limits for education are displayed in figure 2.8. The lower and upper bounds for female eligible subjects are illustrated in red. They vary with the female subject's education which is displayed along the x-axis. For male subjects, the lower and upper bounds are drawn in blue and vary with the subject's educational level displayed along the y-axis. The shaded areas illustrate the 95% confidence intervals. The area between the dotted line and the triangle line each represent the acceptance sets. The correspondence set is where the two acceptance sets overlap.

We observe that the lower bounds for education are increasing in the eligible subject's education for both men and women. Furthermore, the upper bounds are high for both male and female eligible subjects. For example, parents searching on behalf of a female subject with a high school degree would accept men with at least a high school degree and at most a university degree. Parents searching on behalf of a male subject with a high school degree half expect the wife to have at least middle school, half expect at least high school. Some might reject a woman with a university degree, but most would not.


The correspondence set, where the acceptance sets overlap, is bounded by the lower bounds on both sides. This implies that if one side agrees to a match, it is possible that the other side rejects because the proposed person does not have sufficient education. However, a match that is agreeable to one side, would not be rejected by the other side because of too much education.

It is also clear that the matches that are agreeable to both sides are along the 45 degree line: when men and women have the same educational levels. In figure 2.9, we add the information about recently married couples in the general population from the CFPS 2014. The bigger the circle, the higher the frequency of couples at this point. Unfortunately, we do not have information on the limits when the female subjects does not have a high school degree.

For women with a high school degree or higher, 48% are married to a partner who falls within the acceptance set. This increases to 63% when we include the 95% confidence intervals. These results are summarized in table 2.11 in the appendix. In particular, many women with a high school degree are married to a husband who has less education than what we would expect based on the acceptance sets. Parents want their high school educated daughters to marry someone who has a higher educational level. However, as we can see in figure 2.8, most marry someone with a high school degree and several marry a man with only middle school education.

For men with a middle school degree or higher, 77% are married to a woman who falls within the acceptance set. It increases to 91% when we include the confidence intervals. These high numbers are due to parents not having high expectations for their daughter-in-law's education. They usually accept women with the same education or one level lower than their son, which encompasses nearly all of the matches that we observe in the general population.

Overall, 47% of the population fall within the estimated correspondence set, 63% if we include the confidence intervals. Therefore, more than half of the marriages in the general population would have been accepted by the parents that we interviewed. However, parents who search on behalf of an educated daughter might not always have realistic lower educational bounds.




Upper and lower levels that respondents indicated that they would accept as the educational level for the partner of their eligible subject. Responses on behalf of female subjects are displayed on the x-axis. Responses on behalf of male subjects are displayed on the y-axis.

Figure 2.8: Education: Acceptance Sets

The acceptance sets for age are displayed in figure 2.10. Again, the lower and upper limits for female subjects are illustrated in red and vary along the subject's age on the x-axis. The lower and upper limits for male subjects are illustrated in blue and vary along the y-axis.

For age, parents searching on behalf of male and female subjects indicated lower and upper bounds (nobody chooses “no limit”). These bounds are clearly increasing with own age. The accepted age span is usually around 7 year, except for men below the age of 30, where parents have smaller acceptable age spans. Until around the age of 30, the lower limits are bounding the correspondence set: parents of young female subjects want the husband to be older than their daughter. Interestingly, parents of young male subjects do not want them to have a wife that is considerably younger. After the age of around 30, the upper limits are binding. Parents of older female subjects would accept husbands that are around the same age or slightly younger or older, but not one who is considerably older. Parents of older male subjects expect the wife to be younger. The correspondence set gets small toward the end of the 30s.




Upper and lower levels that respondents indicated that they would accept as the educational level for the partner of their eligible subject. Responses on behalf of female subjects are displayed on the x-axis. Responses on behalf of male subjects are displayed on the y-axis. Circles display actual marriages taken from the CFPS 2014.

Figure 2.9: Education: Acceptance Sets with frequency of recent marriages


We overlay the information of marriage outcomes from the general population in figure 2.11. Most people get married in their 20s. Most marriage are in the correspondence set, however, they are bunched around the lower bound of the female subjects. The statistics can be found in table 2.12 in the appendix. Indeed, for women up to 26, as well as for men up to 34, nearly all spouses fall into the estimated acceptance set with confidence intervals (86% to 100%). The numbers are smaller for women between 27 and 34 and men between 35 and 38. The share is also small for the age group over 41, which is very broad but encompasses few observations.

Overall, 40% of couples lie in the correspondence set, 72% in the correspondence set with the confidence intervals. Most marriages that lie outside of the parents' correspondence set involve women between 27 and 34 - whose parents want slightly older sons-in-law (but there are also some who are too old) - or men between 35 and 38 where parents want younger daughter-in-laws.



Upper and lower levels that respondents indicated that they would accept as the age for the partner of their eligible subject. Responses on behalf of female subjects are displayed on the x-axis. Responses on behalf of male subjects are displayed on the y-axis.

Figure 2.10: Acceptance Sets on Age



Upper and lower levels that respondents indicated that they would accept as the age for the partner of their eligible subject. Responses on behalf of female subjects are displayed on the x-axis. Responses on behalf of male subjects are displayed on the y-axis. Circles display actual marriages taken from the CFPS 2014.

Figure 2.11: Acceptance Sets on Age with frequency of recent marriages

2.5.2 Match Simulation

2.5.2.1 Estimation

Can the estimated preference explain current marriage patterns? Marriage is a two-sided matching problem. Outcomes are the combination of preferences on both side. We simulate marriages where the previously estimated preferences determine the demand function and the characteristics of the general population determine the supply.

From the China Family Panel Study (2014) we have information on the educational levels and the age of both spouses for marriages that were formed between 2012 and 2014. We use these observation as supply of husbands and wives. For each woman w (man m), we predict the likelihood of choosing each man m (women w) based on equation 2.3 and the estimated parameters of θ . We denote this likelihood the “selection likelihood”. In the profile taks this is equivalent to the respondent selecting the (male) profile p on behalf of the (female) eligible subject i . Denote Y_{ip} the indicator for denoting “yes” to profile p for subject i . We use θ from three specifications:

- Specification 1 (Baseline): Includes the age and educational difference (profile minus eligible subject):

$$Y_{ip} = \theta_1 \Delta Age_{ip} + \theta_2 \Delta Education_{ip} + \epsilon_{ip} \quad (2.4)$$

- Specification 2: Includes the educational difference interacted with an indicator if the subject has a university degree and the age difference:

$$Y_{ip} = \theta_1 \Delta Age_{ip} + \tilde{\theta}_2 \Delta Education_{ip} ' uni degree_i + \epsilon_{ip} \quad (2.5)$$

- Specification 3: Includes age difference interacted with an indicator if the subject is over 30 and educational difference interacted with an indicator if the subject has a university degree:

$$Y_{ip} = \tilde{\theta}_1 \Delta Age_{ip} ' over 30y_i + \tilde{\theta}_2 \Delta Education_{ip} ' uni degree_i + \epsilon_{ip} \quad (2.6)$$

Specifications 2 and 3 each add an interaction term and thus estimate group-specific preferences. The groups are eligible subjects with and without a university degree and eligible subjects being older and younger than 30. This has the advantage that we can account for characteristic-dependent preferences and potentially get point estimates that are closer to the true preferences. However, since some of the groups do not have many observations in our data set, we also introduce more noise in the estimates. Faced with this trade-off, we run and interpret all three specifications.

The higher the predicted selection likelihood of a woman w for men m (\hat{Y}_{wm}), the higher the utility woman w derives from being matched with men m (see equation 2.1). The predicted selection likelihood thus gives us a cardinal ranking: The man with the highest selection likelihood

is the woman's first choice, the man with the second highest selection likelihood is the woman's second choice and so forth.

The marriage market simulation works as followed: Men propose to the woman that they rank highest based on the predicted selection likelihood. A woman with one proposal stays with the proposer. A woman who received more than one proposal selects the man that she ranks highest among them and rejects the others. In the second round, the rejected men propose to the woman that they attribute the second highest rank. Women then repeat their selection among the proposals. This is repeated until all individuals are matched.

This matching process corresponds to the man-proposing Gale-Shapley mechanism (Gale and Shapley (1962)). The resulting matches are stable. A match is defined as stable if neither of the two spouses can find a better match (Roth and Sotomayor (1992)). However, the matches are not unique. Furthermore, we have several individuals with the same characteristics (age and education) such that they get attributed the same likelihood. The ranking function is programmed to randomly rank individuals with the same likelihood. The matching process is thus repeated 100 times. At the end, we take the average of the match's characteristics over each matching process. In order to verify the robustness of the resulting distribution, we also run the algorithm with women making the proposals and men rejecting and accepting proposals (woman-proposing Gale-Shapley mechanism).

The Gale-Shapley mechanisms forces all women and men to match. There is no option to stay unmarried. This implies that there could be matches that would not form because one side would rather prefer to stay single. To address this issue, we calculate the lowest selection likelihood that is observed in the real outcome data for each specification: For each women w and each man m , we predict the selection likelihood of their actual spouse, based on θ . We determine the minimum of the selection likelihoods for women and men respectively as the minimum thresholds. We can then compute how many individuals have a lower selection likelihood from their simulated match than their sex-specific minimum threshold.

2.5.2.2 Results

Table 2.13 in the appendix displays the results for θ for the three specifications for men and women. These marginal effects are then used to calculate the demand function. As students are always high educated and below 30, we can only use specification 1.

The frequency histogram of the actual educational difference between husband and wife and the simulated educational difference, based on parents' and students' preferences, are displayed in figure 2.12 for the men-proposing Gale-Shapley mechanism. The histograms for the women-proposing mechanism is displayed in figure 2.14 in the appendix. Table 2.5 summarizes the characteristics of

real outcome distribution and the four predicted distributions: the correlation coefficient for the matched couples' levels, the coefficient of variation for the differences between the matched couples' levels, the most common case (mode) and the share of matches with the same level.

Specification 1 based on parents' and students' preferences predict matches that are more assortative on education than what is observed in the actual distribution. More couples would form between husband and wife having exactly the same educational level. Specification 2 and specification 3 based on parents' preferences, however, predict less assortative matches. Specification 2 and 3 based on parents' preferences lead to a higher coefficient of variation in the educational differences between spouses and to a lower correlation between the spouses' educational levels.

The same graphs for the age differences between husband and wife are displayed in figure 2.13 (men-proposing) and 2.15 (women-proposing, in the appendix). Specification 1 based on students preferences predicts again more matches with the same age level. All specification based on parents' preferences shift the peak of the distribution to the right - where husbands are older than their wives. In the actual distribution, the modal point is at the husband and the wife having the same age. In the predicted distributions based on parents' preferences, the modal point is at the husband being 2 years older than the wife in specifications 1 and 2. In specification 3, the mode is at the husband being 3 years older. Again, specifications 2 and 3 based on parents preferences predict distributions with a higher coefficient of variation and lower correlations between the age of spouses.

We calculate how many women and men have a lower selection likelihood for their simulated match that the minimum threshold. In specification 1 and 2 based on parents' preferences and specification 1 based on students' preferences, no woman has a lower selection likelihood than the minimum selection likelihood of real matches. Only in specification 3, 3 women have a lower likelihood than the minimum threshold. For men, it is the reverse: In specification 1 based on parents' preferences, 2 men have a lower selection likelihood that the minimum that men receive from their actual match. In specification 2, 3 men fulfill this criteria and 2 men in specification 1 based on students' preferences. However, in specification 3, all men have a selection likelihood higher than the minimum threshold. As these numbers are small, deleting them from the sample does not change the distributions significantly.

2.5.3 Discussion

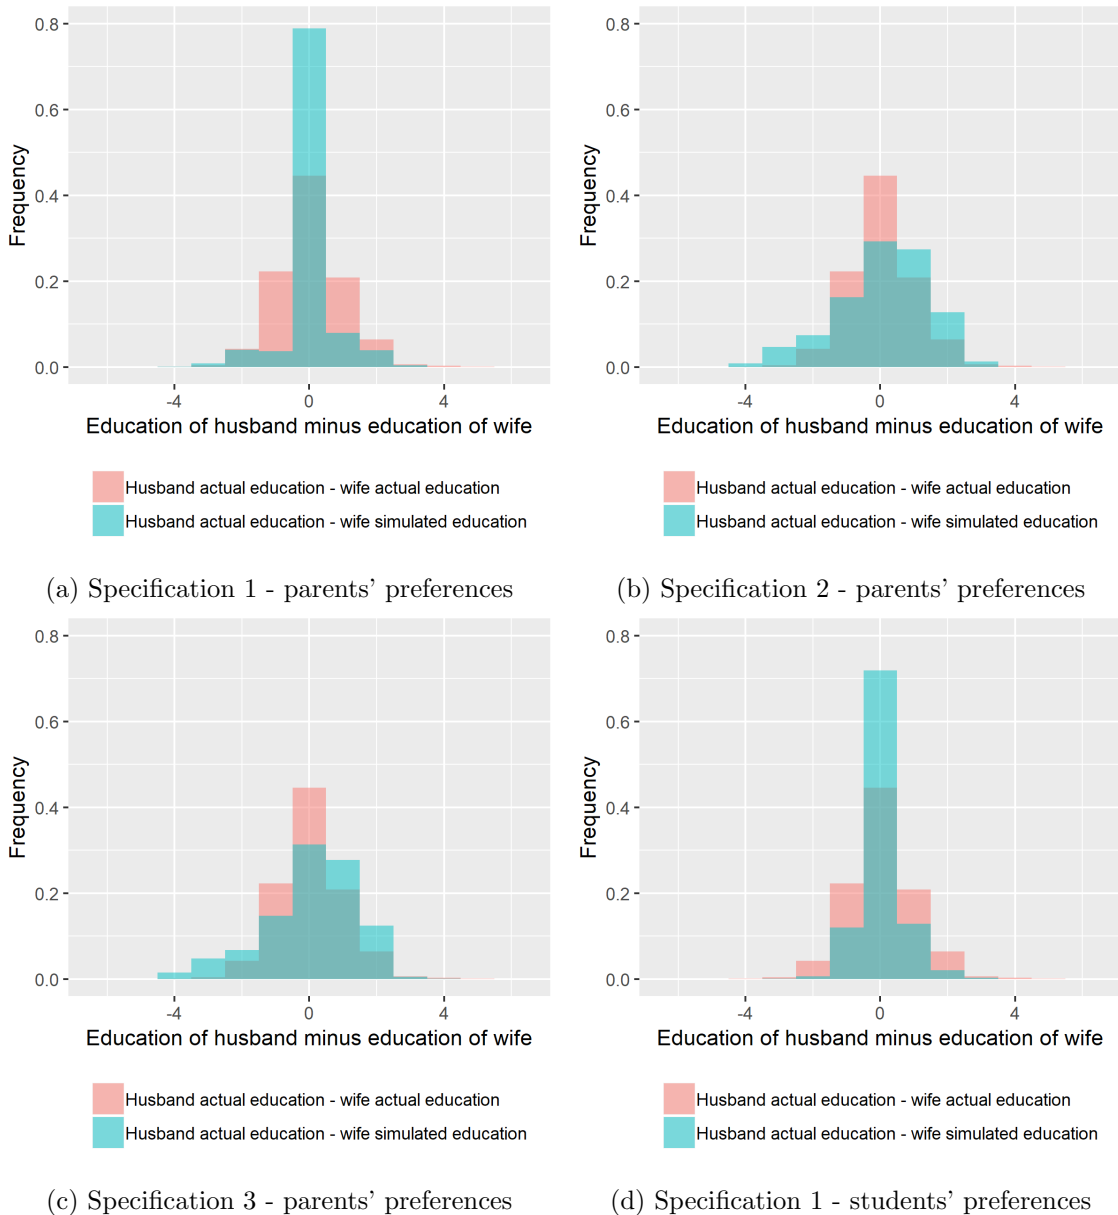
The first observation is that specification 1 with parents' and students' preferences predicts a higher degree of assortativity on education and age. The share of matches with the same educational level is 71% based on students' preferences and 78% based on parents' preferences. In the

¹⁶Pearson's Correlation Coefficient

¹⁷The coefficient of variation of the difference between the spouses.

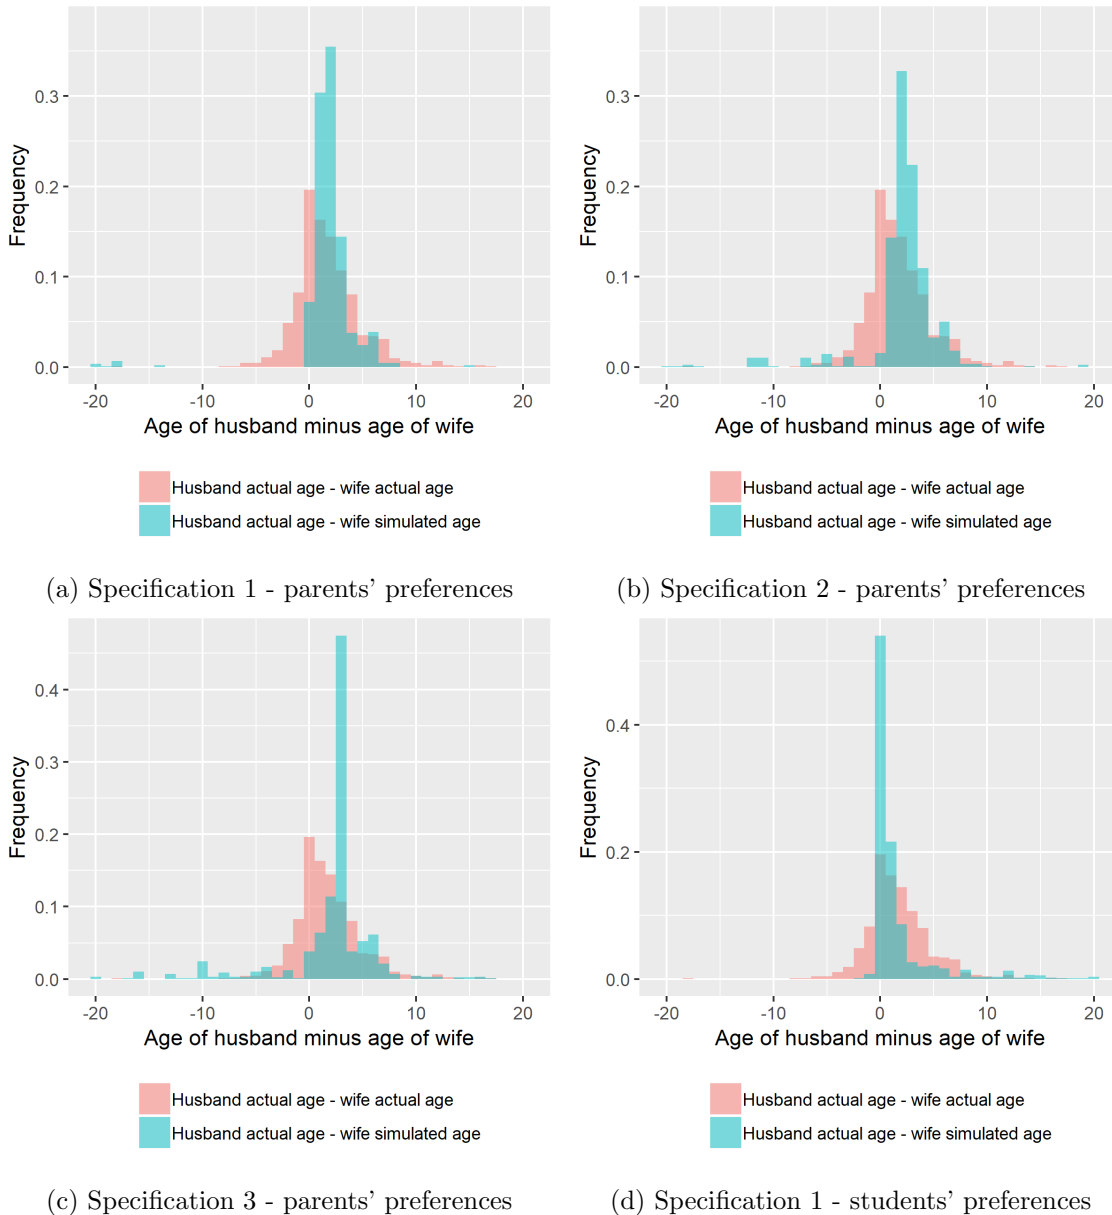
¹⁸Pearson's Correlation Coefficient

¹⁹The coefficient of variation of the difference between the spouses.



Note: Graphs show the actual distribution of the educational difference between husband and wife, taken from the CFPS 2014, including marriages formed between 2012 and 2014. The simulated distribution is based on the real education of the husband and the simulated education of the wife. Preferences taken from table 2.13. Uses the men-proposing Gale-Shapley Algorithm.

Figure 2.12: Education: Comparing actual educational differences between husband and wife and educational differences in simulated matches.



Note: Graphs show the actual distribution of the age difference between husband and wife, taken from the CFPS 2014, including marriages formed between 2012 and 2014. The simulated distribution is based on the real age of the husband and the simulated age of the wife. Preferences taken from table 2.13. Uses the men-proposing Gale-Shapley Algorithm.

Figure 2.13: Age: Comparing actual educational differences between husband and wife and educational differences in simulated matches.

	Real distribution	Spe. 1 Parents	Spe. 2 Parents	Spe.3 Parents	Spe. 1 Students
Education					
Average difference btw. spouses	0.05	0.05	0.05	0.05	0.05
Correlation between spouses' levels ¹⁶	0.56	0.44	0.03	0.01	0.56
Coefficient of Variation ¹⁷	20.71	14.78	26.95	27.82	12.01
Most common case (mode)	Same level	Same level	Same level	Same level	Same level
Share same level	0.45	0.78	0.29	0.31	0.71
Age					
Average difference btw. spouses	1.68	1.68	1.68	1.68	1.68
Correlation between spouses' levels ¹⁸	0.87	0.81	0.61	0.64	0.83
Coefficient of Variation ¹⁹	1.91	2.31	3.33	3.16	2.18
Most common case (mode)	Same age	Husband 2 years older	Husband 2 years older	Husband 3 years older	Same age
Share same age +/- 2 years	0.63	0.72	0.48	0.22	0.84

Table 2.5: Comparing real marriage distribution with predicted distributions

real distribution, the share is only 45%. The share of matches with a small age gap (up to 2 years) is 72% based on parents' preferences and 84% based on students' preferences, compared to 63% in the real distribution. This higher degree of assortativity can be explained by the lack of search frictions in the simulation. Search frictions in the marriage market can include search costs or limits to whom you meet. However, the non-representativeness of our sample can also be an explanation. The preferences in the general population might be more heterogeneous than the preferences in our sample.

Specification 2 and 3 lead to higher variation in the differences between the levels of spouses for both education and age. These specifications have group-specific estimates. These group-specific estimates, though possibly closer to the true parameters, are also noisier. This is particularly true for the groups that are under-represented in our survey: those without a university degree and below 30. Yet, these are the groups that are more represented in the general marriage population. The noisily measured estimates can explain the low correlation levels and the high coefficients of variation.

The distributions based on parents' and students' preferences display one clear difference: Students' preferences predict the mode of the distribution to be where spouses have the same age. Parents' preferences predict the modal point to be at an age difference between 2 and 3 years. In this case, the students' preferences are more reflected in the actual outcome data. This could be because the eligible subjects on whose behalf parents are searching are on average older than the average marriage age. Parents with an older son might prefer a higher age gap for fertility reasons - if they believe that the likelihood of having a grand-child is higher if the daughter-in-law is below 30 rather than in her 30s. Yet, accounting for different preferences for eligible subjects above and below 30 only increases the age gap in the simulations.

We know from section 2.5.1 that when the eligible subjects are in the early or mid-20s, matches between individuals with the same age are acceptable for parents. It seems that a lot of actual matches are preferable for students and acceptable for parents. This hints at an interesting dynamic between parents and their adult children when it comes to spouse choice. This encourages further research.

When looking at the acceptance sets on age, over 80% of actual husband and wives lie within the acceptance set of parents. More than 90% of wives fall within the acceptance sets for parents. Yet, nearly 40% of husbands are outside of the acceptance set of parents. Parents seem to have excessively high expectations on the educational level of their son-in-law. Again, this could be due to our non-representative sample. The overall population might have wider acceptance sets on education. It can also be due to parents having unrealistic expectations about the marriage outcome of their daughter. To form accurate expectations, parents need to know the distribution

in types and the preferences rankings for both sides of the market. This implies having information about the educational distribution of unmarried men and women and their preferences on education. This might be difficult and costly to obtain. Finally, respondents could also answer this question strategically. The acceptance sets could give us limits to what they aspire or expect rather than what they would accept if they had no other choice.

2.6 Conclusion

This chapter investigates marriage sorting and spousal preferences in China. It discusses not only the preferences of the individual that is single, but also those of their their parents and other relatives. Marriage is an important institution in China: only 1% of individuals between 40 and 65 have never been married (China Family Panel Study 2014). Parents with children above the average marriage age might pressure their children or actively search for a suitor themselves. We interview parents that fall at least into the latter category.

The parents we interviewed see themselves as agents for their children. They often state that their child is too busy or too shy to search for themselves. They might also search in parallel with their children. They also rank altruistic search motives as more important than more selfish motives.

Yet, we find that parents' preferences and preferences of students do not always overlap and would not lead to the same marriage patterns. Parents seem to have a preference for an age gap that students do not. This could be due to age preferences changing with the age of the eligible subject. Parents mostly search on behalf of an eligible subject that are older than students. Yet, it is also possible that there is a difference between generations. If the latter is true, where would this come from? One possibility is that the younger generation is more in contact with other individuals of the opposite sex. This could be due to the increase in time spent on education. Higher secondary and tertiary education is becoming more and more common in China. It implies that individuals meet others with the same age (and education) during this time where many couples form. These meeting opportunities, seeing couples form around them, can shape their preferences. While matches between spouses with the same age are not necessarily the equilibrium outcome based on parents' preferences, they still accept those marriages. This motivates further research into the dynamics of bargaining over spousal choice.

Both parents' and students' preferences predict that most matches form between spouses with the same level of education. Yet, many parents searching on behalf of a female subject accept only sons-in-law with a higher educational level. Even if those parents in the end settle for a son-in-law with the same education, trends in the educational distribution might be worrisome for them. Since 2016, more women are enrolled in tertiary education than men.²⁰ This suggests that in the future,

²⁰According to the UNESCO Institute for Statistics, the gross tertiary female school enrollment rate was 53%


more women might have to settle for a husband with a lower level of education or stay single.²¹

2.7 Appendix

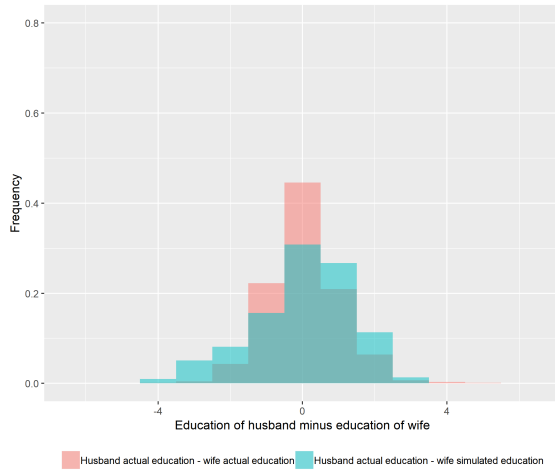
2.7.1 Figures

in 2016 and 56% in 2017. Accessed on the World Bank website <https://data.worldbank.org/indicator/SE.TER.ENRR.FE?locations=CN> in February 2019.

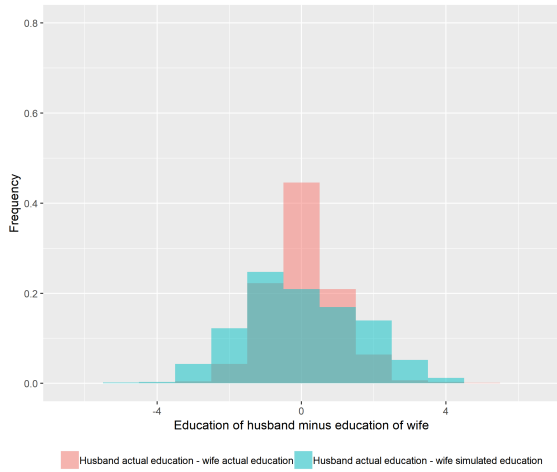
²¹This chapter is a part of the collaboration of Eva Raiber, Weiwei Ren, Jeanne Bovet, Paul Seabright and Charlotte Wang. The initiative for the project was started by Eva Raiber, Paul Seabright and Jeanne Bovet. Charlotte Wang and Weiwei Ren successfully organised and coordinated the cooperation with the Yunnan Normal University. The questionnaire and the profile task used for this chapter was proposed by Eva Raiber, with important contributions from all team members, and translated by Weiwei Ren, Charlotte Wang and the Yunnan Normal University student team. Weiwei Ren, with the support of Charlotte Wang, recruited the student team and trained them with the help of Eva Raiber and Jeanne Bovet. Weiwei Ren organised and supervised data collection. Data analysis and write up of the chapter was done by Eva Raiber with inputs from Paul Seabright and comments by the other team members.




(a) Specification 1 - parents' preferences



(b) Specification 2 - parents' preferences




(c) Specification 3 - parents' preferences



(d) Specification 1 - students' preferences


Note: Graphs show the actual distribution of the educational difference between husband and wife, taken from the CFPS 2014, including marriages formed between 2012 and 2014. The simulated distribution is based on the real education of the husband and the simulated education of the wife.

Figure 2.14: Education: women-proposing Gale-Shapley algorithm



(a) Specification 1 - parents' preferences

(b) Specification 2 - parents' preferences



(c) Specification 3 - parents' preferences

(d) Specification 1 - students' preferences

Note: Graphs show the actual distribution of the age difference between husband and wife, taken from the CFPS 2014, including marriages formed between 2012 and 2014. The simulated distribution is based on the real age of the husband and the simulated age of the wife.

Figure 2.15: Age: women-proposing Gale-Shapley mechanism

2.7.2 Tables

Husband: Educational Level by wife: educational level

Husband: Educational Level	Wife: Educational Level					
	Illiterate Semi-literate	Primary school	Junior HS	Senior HS	College	Bachelor's degree
Illiterate/Semi-literate	13	8	4	0	1	0
Primary school	5	32	56	17	2	0
Junior high school	11	39	167	69	11	2
Senior high school	1	12	65	71	36	9
College	0	4	21	38	61	28
Bachelor's degree	2	3	1	16	39	50
Graduate degree	0	0	0	0	2	7
Total	32	98	314	211	152	96

Note: Includes those that married

Table 2.6: Educational Levels of husband and wife

Variable	Mean	Std. Dev.	N
Respondent: Female	0.678	0.468	549
Respondent: Mother	0.481	0.5	549
Respondent: Father	0.158	0.366	549
Respondent: Uncle	0.016	0.127	549
Respondent: Aunt	0.053	0.224	549
Respondent: Married	0.915	0.28	410
Respondent: Widowed	0.049	0.216	410
Respondent: Divorced	0.016	0.127	549
Respondent: Retired	0.84	0.367	412
Respondent: Working	0.189	0.392	412

Note: Source QSAMPY 2016

Table 2.7: Summary statistics on respondents

²²Mothers and fathers, excluding uncles and aunts and other relatives

	All ²²	Searching Husband		Searching Wife	
		Mother	Father	Mother	Father
Eligible subject having someone who helps him/her with household chores .	52%	41%	68%	62%	57%
Eligible subject having someone to take care of him/her in time of need (sickness/old-age).	86%	85%	79%	90 %	81%
Eligible subject having someone he/she feels affection for.	96%	95%	97 %	97 %	98%
Having someone to care for you (and your spouse) when you are older.	51%	51%	50%	56%	43%
Having grandchildren	63%	60%	57%	69%	68%
N	351	170	34	94	53

Source: QSAMPY (2016)

Table 2.8: Share of respondent indicating reason for search as important or very important.

Variable	Mean
Age	33.063 (5.292)
Child: Female	0.58
Completed at least (professional) high school	0.98
Completed university degree	0.6
father	0.212
mother	0.644
Child: Has been married	0.08
N	410

Note: Source QSAMPY 2016

Table 2.9: Summary statistics on eligible subjects on whose behalf respondents are searching

Dependent variable: Indicator of preferring profile 1.

	Women	Men	Women without univ. degree	Men with univ. degree	Women with univ. degree	Men
	(1)	(2)	(3)	(4)	(5)	(6)
Age difference (1 minus 2)	0.000261 (0.00112)	-0.0137*** (0.00293)	0.00270 (0.00186)	-0.0106*** (0.00403)	-0.00479 (0.00295)	-0.0151*** (0.00414)
Diff.education (1 minus 2)	0.0975*** (0.0116)	0.0261* (0.0155)	0.0817*** (0.0221)	-0.00292 (0.0206)	0.104*** (0.0129)	0.0756*** (0.0221)
Income difference (1 minus 2)	0.0104*** (0.00210)	0.00233 (0.00293)	0.0163*** (0.00397)	-0.00266 (0.00426)	0.00851*** (0.00241)	0.00699* (0.00370)
Real estate diff.(1-2)	0.0510** (0.0223)	0.0320 (0.0264)	0.0382 (0.0413)	0.00563 (0.0344)	0.0532** (0.0262)	0.0656* (0.0381)
Han ethnicity diff. (1-2)	0.0452** (0.0214)	0.0240 (0.0267)	0.127*** (0.0423)	0.00566 (0.0338)	0.0132 (0.0236)	0.0443 (0.0420)
Observations	932	635	275	347	657	288

Note: Logit regression, coefficients are average marginal effects, standard error in parenthesis are clustered pm the individual level. Source: QSAMPY 2016

Table 2.10: Estimation using hypothetical choice data and the question which of the two profiles the respondent prefers for the eligible subject.

	Educational Level	N	Average Lower bound	Average Upper bound	Share partners in AC	Share in CS	Share partners in AC + CI	Share in CS + CI
Women	High School (Level 4)	204	4.09	6.06	0.23		0.58	
	College (Level 5)	161	4.8	6.63	0.7		0.7	
	University (Level 6)	107	5.37	6.59	0.63		0.63	
	Total	472	4.62	6.37	0.48		0.63	
Men	Middle School (Level 3)	301	2.29	5.71	0.83		0.96	
	High School (Level 4)	193	3.45	5.78	0.55		0.93	
	College (Level 5)	150	3.82	6.06	0.87		0.87	
	University (Level 6)	132	4.82	6.4	0.85		0.85	
	Total	776	3.3	5.91	0.77		0.91	
						0.47		0.63

Note: Average age, lower and upper bounds according to age categories. Column 5: Share of partners that fall into the acceptance set. Column 7: Share of partners that fall into the acceptance set plus the 95% confidence intervals. Column 6: Share of matches that are in the correspondence set. Column 8: Share of women that are in the correspondence set plus the 95% correspondence intervals.

Table 2.11: Lower and upper acceptance bounds on education according to parents

	Age category	N	Average Lower bound	Average Upper bound	Share partners in AC	Share in in CS	Share partners in AC + CI	Share in CS + CI
Women	<23	209	22.5	31.75	0.66		1	
	23-26	414	24.65	31.24	0.71		0.86	
	27-30	179	28.13	34.18	0.48		0.64	
	31-34	46	31.93	38.06	0.48		0.48	
	35-38	21	35.61	42.17	0.81		0.81	
	39-41	6	39.17	45.17	0.5		0.5	
	>41	34	45.73	54.64	0.29		0.68	
	Total	909	26.35	33.5	0.62		0.82	
Men	<23	92	21.33	24.67	0.43		0.97	
	23-26	364	21.33	25.8	0.65		0.86	
	27-30	285	23.41	28.53	0.72		0.88	
	31-34	82	25.1	31.92	0.61		0.9	
	35-38	35	27.89	33.91	0.46		0.49	
	39-41	15	29.78	36.89	0.47		0.87	
	>41	36	34.25	43.19	0.25		0.47	
	Total	909	23.23	28.28	0.62		0.85	
					0.4		0.72	

Note: Average age, lower and upper bounds according to age categories. Column 5: Share of partners that fall into the acceptance set. Column 7: Share of partners that fall into the acceptance set plus the 95% confidence intervals. Column 6: Share of matches that are in the correspondence set. Column 8: Share of women that are in the correspondence set plus the 95% correspondence intervals.

Table 2.12: Lower and upper acceptance bounds on age according to parents

	<i>Dependent variable: Indicator: Select Profile</i>							
	Female Subject			Male Subject			Female Students	Male Students
	Spe. 1 (1)	Spe. 2 (2)	Spe. 3 (3)	Spe. 1 (4)	Spe. 2 (5)	Spe. 3 (6)	Spe.1 (7)	Spe 1 (8)
Age Difference (+)	-0.129*** (0.021)	-0.130*** (0.021)	0.051 (0.073)	-0.166*** (0.039)	-0.171*** (0.039)	-0.112 (0.136)	-0.0001 (0.046)	-0.099*** (0.033)
Age Difference (+) Sqrd	0.003*** (0.001)	0.003*** (0.001)	-0.008 (0.005)	0.004** (0.002)	0.004** (0.002)	-0.006 (0.013)	-0.004 (0.003)	0.002 (0.002)
Age Difference (-)	-0.388*** (0.054)	-0.388*** (0.054)	-0.102 (0.235)	0.186*** (0.049)	0.188*** (0.049)	0.411 (0.282)	-0.263 (0.176)	-0.143* (0.084)
Age Difference (-) Sqrd	0.023*** (0.005)	0.023*** (0.005)	-0.023 (0.045)	-0.015*** (0.004)	-0.015*** (0.004)	-0.075 (0.052)	-0.012 (0.032)	-0.001 (0.012)
Education Diff (+)	0.126 (0.101)	0.017 (0.206)	-0.002 (0.207)	-0.132 (0.109)	-0.115 (0.153)	-0.140 (0.155)	0.255** (0.107)	0.142* (0.080)
Education Diff (-)	-0.557*** (0.079)	-0.532* (0.312)	-0.551* (0.310)	-0.276*** (0.089)	0.074 (0.178)	0.051 (0.180)	-0.448*** (0.100)	-0.524*** (0.074)
Indicator: Educated		0.189 (0.230)	0.194 (0.231)		0.096 (0.199)	0.065 (0.201)		
Education Diff(+) X Ind.: Educated		0.220 (0.239)	0.236 (0.240)		-0.002 (0.231)	0.048 (0.234)		
Education Diff(-) X Ind.: Educated		-0.039 (0.323)	-0.019 (0.321)		-0.442** (0.207)	-0.408* (0.209)		
Indicator: Over 30			0.411 (0.257)			-0.747** (0.322)		
Age Difference(+) X Ind.: Over 30			-0.211*** (0.078)			-0.041 (0.147)		
Age Diff. (+) Sqrd X Ind.: Over 30			0.011** (0.005)			0.010 (0.013)		
Age Difference(-) X Ind.: Over 30			-0.298 (0.243)			-0.151 (0.288)		
Age Diff. (-) Sqrd X Ind.: Over 30			0.047 (0.045)			0.057 (0.052)		
Observations	1,845	1,845	1,845	1,266	1,266	1,266	1,264	2,238

Source: QSAMPY 2016. Logit regression with non clustered standard errors. Education variable more broadly defined than in previous tables to fit to the variable used in the CFPS 2014. Indicator for educated is 1 if the subject has college or university education. *p<0.1; **p<0.05; ***p<0.01

Table 2.13: Estimations to simulate marriage outcomes

Chapter 3

God insures those who pay? Formal insurance and religious offerings in Ghana¹

EMMANUELLE AURIOL, JULIE LASSÉBIE, AMMA PANIN, EVA RAIBER,
PAUL SEABRIGHT

¹ This chapter is a revised version of the TSE Working Paper “God insures those who pay? Formal insurance and religious offerings” (Auriol et al. (2017)). This project is the result of a collaboration between Central University in Accra, Ghana, and the Institute for Advanced Study in Toulouse - Toulouse School of Economics. We are immensely grateful to Anthony Panin for coordinating the collaboration between CU and IAST-TSE. We also thank our field assistants Bernand Boakye and Davis Adu Larbi. We would like to thank Astrid Hopfensitz, Michael Kosfeld, Michael McBride and Josepa Miquel-Florensa for useful comments. We also thank participants at IAST and TSE seminars, CIFAR meeting in Paris, and ASREC meetings for their valuable feedback. Support through the ANR - Labex IAST is gratefully acknowledged, and also from the ERC COGNITION Grant Agreement 249429.

ABSTRACT

This chapter provides experimental support for the hypothesis that insurance can be a motive for religious donations by members of a Pentecostal church in Ghana. We randomize enrolment into a commercial funeral insurance policy, then church members allocate money between themselves and a set of religious goods in a series of dictator games with significant stakes. Members enrolled in insurance give significantly less money to their own church compared a control group of members that only receive information about the insurance. Enrolment also reduces giving towards other spiritual goods. We set up a model exploring different channels of religiously based insurance. The implications of the model and the results from the dictator games suggest that adherents perceive the church as a source of insurance and that this insurance is derived from beliefs in an interventionist God. Survey results suggest that material insurance from the church community is also important and we hypothesize that these two insurance channels exist in parallel.

3.1 Introduction

Do religious believers give money to their churches in the hope of receiving insurance against economic shocks? And if so, is this because they expect the church to look after them when shocks occur, or God to look after them by making such shocks less likely to happen? Recent studies have emphasized the important economic functions that religious institutions fulfil across the world, of which the provision of informal insurance is one that has recently gained attention (Chen (2010), Ager and Ciccone (2016)). Such insurance is likely to be particularly important in settings with weak formal institutions and where individuals are exposed to a variety of uninsured risks. In this chapter we report a study examining the effect of providing formal market-based insurance on the demand for informal church-based insurance in Accra, Ghana.

We conduct a lab-in-the-field experiment with members of a Pentecostal church in Accra to test whether insurance is one motive behind religious participation. We do not claim it is the only or even the most important motive. We perform our experiment by randomly assigning free enrollment into a formal, commercially available funeral insurance policy and measuring how this affects willingness to contribute money in a dictator game to the church and two other charitable recipients.

The additional recipients - a secular charity and a nationwide prayer event - provide a means of differentiating between two possible mechanisms by which members might expect insurance to work. The first mechanism, consistent with a club good interpretation of church involvement, is one in which the church provides financial support to those who suffer economic hardship provided they have previously given money to the church. The second mechanism is one in which members expect God to intervene to reduce the risk of adverse economic shocks provided the individuals concerned have demonstrated spiritually meritorious actions. Giving money to a secular charity or to a nationwide prayer event might be considered meritorious actions but they do not involve a contribution to the church, and so would be consistent with belief in the second mechanism but not with belief in the first.

Urban Ghana is a particularly interesting setting in which to study interactions between religious participation and insurance. First of all, as we describe in section 2.2.1, Pentecostalist churches are massively present in Ghana and many members of the population give very significant shares of their income to their church. This setting therefore allows us to examine the motives for costly participation in activities that are also present in many other countries in the world. In our model we envisage multiple motives for participation - both the consumption of spiritual activities as such and the purchase of insurance against various shocks, and our survey evidence indicates that both of these motives matter. Secondly, in our main sample of church members we find that only 31% of all those interviewed participate in the National Health Insurance Scheme (NHI), Ghana's public health insurance program. Only 19% indicate that they hold any other types

of insurance (see Table 3.8). Relatively low rates of insurance mean that enrollment in a formal policy is likely to be a meaningful and significant treatment for many of our participants. Overall, therefore, our results are likely to be relevant for other settings characterized by low penetration of formal insurance and the importance of religious institutions, including other Sub-Saharan African countries but also other parts of the world such as Latin America.

We find that enrollment in the formal insurance policy causes church members to give less money to the church in a dictator game, compared to members that only receive information about the insurance. Interestingly, we find that formal insurance also causes church members to give less to the other recipients who are not directly linked to the church, but are associated with church teachings on “good behaviour”. Furthermore, we find that church members who receive only information about the funeral insurance but are not enrolled, increase their giving to the church and other charitable recipients as compared to participants who are not informed about the content of the insurance. This finding is interesting since many of the church’s communications place a heavy emphasis on the ever-present risk of death, and our results suggest that information which increases the salience of this risk is highly effective at increasing individual willingness to donate.

We set up a simple model to illustrate the conditions under which we can interpret these results as evidence for insurance. In the model, a church member derives utility from secular and religious consumption. We first show that if religious consumption does not have an insurance motivation, a reduction in perceived losses (through our enrolment treatment) should increase the amount allocated to religious consumption, via an income effect. However, if religious consumption does have an insurance motivation, and if the insurance is perceived as sufficiently effective, the substitution effect of a reduction in perceived losses will outweigh the income effect and lead to a reduction in the amount allocated to religious consumption. This is what we find in our experiment.

This insurance motive can work either through a reduction in perceived losses from negative shocks because of transfers from other church members or from the church organisation, or through a reduction in perceived probabilities of adverse shocks because of belief in a divine response to religious giving. Our experiment finds strong evidence for the latter mechanism: church members change donations to the church as well as to the non-church recipients in a manner that is consistent with interpreting giving as a demand for divine protection. Indeed, church members who receive only information about funeral insurance react to the increased salience of a potential loss with an increase in donations to the church and other non-church recipients. Provision of formal insurance decreases the need for this type of risk-coping mechanism, and we find that insured church members decrease charitable giving compared to those that are primed with risk.

Finally, we find that subjects who give substantially more than the average, or who participate in the game during intensive “revival weeks” – events in which churchgoers donate much more than

in an average week – demonstrate treatment effects that are of the opposite sign to the average for the sample. This is consistent with the interpretation that exogenous factors that increase religious giving, thereby reducing the marginal insurance utility from religious consumption, can reduce the probability that the substitution effect outweighs the income effect.

Although our model is an instance of standard micro-economic analysis as applied to the allocation of resources between secular and religious consumption, our empirical study is methodologically innovative in two main ways. First, we have provided an experimental intervention that directly affects the demand for religious insurance by providing an institutional substitute. Secondly, we have sought to distinguish carefully between different channels through which religious insurance might work - notably, through influencing the behavior of other members of the religious community as opposed to influencing the probabilities of adverse events themselves that are believed by subjects to be determined by God. To this end, our study experimentally links the religious practice of church giving to an insurance motive, and we find evidence that subjects believe their donations can induce God to intervene causally in the world. A Pentecostal church is a particularly good setting to test this because participation costs and expected outcomes are made explicit, and “giving to God” has a clear doctrinal mandate.

3.1.1 Links with the literature

The macro-economic literature, in the wake of [Weber \(1905\)](#), has long recognized the potential importance of culture, and especially religion, for economic growth. However, most empirical studies in this literature face a challenge in surmounting endogeneity problems, and it is generally hard to rule out the possibility that confounding factors explain both the religiosity of a population and the growth of its economy. For instance, [McCleary and Barro \(2006a\)](#) find in their instrumental variable model that higher GDP per capita causes a reduction in average religiosity. They also find that the relationship between economic development and religiosity depends on the particular dimension of development. For example, education has different effects than urbanization ([McCleary and Barro \(2006b\)](#)). Thus, even when it is possible to estimate the magnitude of aggregate causal effects, such large-scale studies have difficulty pinpointing the mechanisms involved. Experimental methods are therefore helpful both in establishing causality and in identifying the likely mechanisms. This chapter aims at understanding a specific economic function of religious organizations: their role as informal insurers.

In examining religious organisations as insurers, our work follows that of [Chen \(2010\)](#), who finds that religious intensity increased with the need for ex-post insurance after financial shocks in Indonesia, and [Ager and Ciccone \(2016\)](#), who find a relationship between higher rainfall risk and religious participation. However, each study is compatible with a number of possible causal mechanisms, which our experiment is designed to elucidate. In particular, while the findings of

Chen (2010) suggest an insurance motive for religious participation, they are also compatible with other explanations for religious participation, such as that religion provides comfort in distress after shocks strike. The findings of Ager and Ciccone (2016) are compatible with interpretations where other characteristics associated with rainfall risk are also the proximate cause of religious membership (such as the amount and type of education). Our experiment allows us to allocate insurance randomly to subjects to test the insurance mechanisms more directly.

In other studies, religious participation has also been shown to provide partial insurance against fluctuations in consumption and well being (Dehejia et al. (2007)). Other evidence from cross-country surveys and historical evidence from the Great Depression demonstrate a degree of substitution between access to social welfare and religious participation (Gruber and Hungerman (2007); Scheve and Stasavage (2006)). We add experimental evidence to this literature, causally demonstrating that access to secular insurance can reduce religious involvement. More importantly, the experimental approach allows us to test for the mechanism at hand and therefore acts as a complement to previous research done with observational data.

Our finding that spiritual insurance is an economically important reaction to risk of those that are not formally insured has important implications for the interpretation of the previously mentioned studies. Indeed, the previous literature on the interplay between religion and insurance is not able to distinguish between those two channels and might overstate the role of religious institutions as a provider of community-based financial insurance and understate the importance of a spiritual or psychological response to risk that is reflected in increased religiosity. This problem of identification is vindicated by some recent studies showing that beliefs in divine intervention in the daily lives of individuals can be an important determinant of real and costly social decisions (see Hadnes and Schumacher (2012), Gershman (2016) and Nunn and de la Sierra (2017)).²

Our experiment contributes to the microinsurance literature by studying in a controlled environment how religious participation might affect demand for formal insurance. In recent years, microinsurance policies have been proposed and tested as poverty alleviation tools with limited success (Cole et al. (2013), Giesbert et al. (2011), Giesbert and Steiner (2015), Karlan et al. (2014)). Consistently across studies, take-up has been lower than expected and this has been attributed to a variety of factors including liquidity constraints (Cole et al. (2013)), limited attention (Zwane et al. (2011)), trust in the insurance mechanism (Karlan et al. (2014)), and the existence of informal insurance substitutes (Mobarak and Rosenzweig (2013)). Our results suggest that religious institutions, which are an important instance of informal networks in developing countries, already provide global spiritual insurance to their members (although we do not claim that this is the only

²Gershman (2016) demonstrates a negative association between trust and the presence of witchcraft beliefs, while Hadnes and Schumacher (2012) experimentally find that priming voodoo beliefs increases trust and trustworthiness. Nunn and de la Sierra (2017) document the prevalence of beliefs in protective spells in South Congo, arguing that these beliefs helped village residents coordinate a stronger resistance against enemy fighters.

motive for church participation), and this may partly explain the puzzle of the limited success of specific formal insurance in some contexts. Since our survey results suggest that churches are also important sources of financial assistance to their members, we hypothesize that an informal, community-based material insurance exists in parallel, reducing further the need for formal insurance.

In the following section, we describe the religious context of our study. In section 3, we present the experimental design. In section 4, we use a simple model to derive our experimental hypotheses to identify an insurance effect and to distinguish between community-based insurance and spiritual insurance. In section 5, we discuss the experimental results and conclude in section 6.

3.2 Context

3.2.1 Pentecostalism

Pentecostalism represents one of the fastest-growing segments of global Christianity. Approximately one-quarter of the world’s two billion Christians are members of churches that can be classified as Pentecostal or Charismatic. These are related movements of Protestant Christianity that emphasize the work of a Holy Spirit and claim that spiritual gifts, such as prophecy, divine healing and speaking in tongues are expected to be present in the lives of believers ([Pew Research Center \(2006\)](#)). Although Pentecostal and Charismatic movements differ in some aspects, these aspects are marginal for our study, and we will use the term Pentecostal broadly. This makes sense especially in Ghana and across Africa, where popular speech hardly draws a distinction between the two ([Okyerefo \(2011\)](#); [Botha \(2007\)](#)).

The striking growth of this movement of global Christianity has largely taken place over the last three decades ([Botha \(2007\)](#); [Thelen \(2017\)](#)). The speed with which Pentecostal congregations have grown is particularly remarkable as most typical definitions of the movements only date their beginnings to the early 1900s. This marked change in the composition of global Christianity has been so significant that a number of commentators refer to the “Pentecostal explosion” as a “new Reformation” ([Thelen \(2017\)](#); [Jenkins \(2011\)](#)).³ Sub-Saharan Africa has been an important centre of this change. Estimates from 2015 suggest that almost 40% of the continent’s Christians identify as Pentecostal or Charismatic.⁴ This reflects the wider shift in the Pentecostal revival from

³Comparing the reformation of 16th century Europe with the “new reformation” of Pentecostal-charismatic Christianity, theologian Harvey Cox states that “today, Christianity is living through a reformation that will prove to be even more basic and more sweeping than the one that shook Europe during the sixteenth century . . . That earlier reformation . . . was confined to one small corner of the globe. The current reformation, however, is an earth-circling one. The present reformation is shaking the foundations more dramatically than its sixteenth-century predecessor, and its results will be more far-reaching and radical.” [Cox \(2011\)](#)

⁴Todd M. Johnson and Gina A. Zurlo, eds. World Christian Database (Leiden/Boston: Brill, 2007). <http://worldchristiandatabase.org>

the Global North to the Global South. Although many early Pentecostal churches in sub-Saharan Africa were planted and funded by American missionaries, much of the dramatic growth that has taken place since the 1980s has been fuelled by indigenous developments in theology, teaching and practice. Many scholars argue that preexisting belief systems and economic conditions made sub-Saharan Africa particularly receptive to Pentecostal teachings (see for instance [Ukah \(2005\)](#)).

Participants in our experiment were recruited from different branches of the Assemblies of God (AoG), one of the oldest and most established Pentecostal denominations. Founded in 1914 in the United States of America, the denomination now has an international reach that covers approximately 360 000 churches in more than 250 countries ([Assemblies of God World Missions \(2017\)](#)). National AoG bodies are largely autonomous but are united by some shared beliefs and a common history. Taken together, Assemblies of God fellowships form the largest Pentecostal denomination and the sixth largest religious grouping in the world. ⁵

About a third of the worldwide Assemblies of God adherents are found in sub-Saharan Africa. This represents approximately twenty-two million adherents ([Assemblies of God World Missions \(2017\)](#)), about 2.5 million of whom are in Ghana⁶. In Ghana they represent just under 10% of the national population, attracting members from all social strata. Table 3.4 in the appendix presents some descriptive statistics of the six church branches where we recruited participants for our study. We approached these churches using a snowball sampling strategy. The church branches are heterogeneous in terms of age, size, members' characteristics, and geographical locations in the city. They represent the great diversity of Assemblies of God branches that can be found in Accra.

Many of these teachings that have taken root and been developed in sub-Saharan churches are shared by Pentecostal believers across the world in the transnational exchange of ideas and resources that characterises Pentecostalism today ([Gifford \(2004\)](#); [Kalu \(2008\)](#)). There is, of course, considerable heterogeneity in beliefs and practices within individuals in a given church, across branches within a single denomination, and across denominations of the broader Pentecostal movement. However, the key features of Pentecostal belief and practice which are important for the interpretation of our study design are widely shared. By focusing recruitment from a popular and established denomination, we are able to draw out insights that should be applicable to this large and growing movement, and to other religious practices that share certain features of these beliefs.

⁵www.adherents.com

⁶Todd M. Johnson and Gina A. Zurlo, eds. World Christian Database (Leiden/Boston: Brill, 2007). <http://worldchristiandatabase.org>

3.2.2 The role of the church in Pentecostalism

In terms of practice, the church is an essential part of life for Pentecostal adherents. They go to church more regularly than other Christians and perform other religious practices more frequently. For example, compared to the general population, Pentecostals pray and read the Bible more often, and more frequently watch or listen to religious programs on television and radio. They are also more likely to share their beliefs with others to spread their faith (Pew Research Center (2006)).

Members also enjoy non-spiritual benefits from their church. An essential function of Pentecostal churches in Ghana, in particular in urban areas, is to offer a place for social gathering. Table 3.5 in the appendix shows descriptive statistics of selected answers from our study questionnaire about the social role of churches. For instance, 40% of our study participants declare that they have found their spouse (or are most likely to find their future spouse if not already married) in church. This does not mean that for them, the church is just or even mainly a place to meet potential marriage partners. On the contrary, precisely its appeal in this dimension consists in providing access to marriage partners who are there for other, at least partly spiritual or ethical motives.

Some church members even report that they favour church members as business partners (48% of church members in our study). More fundamentally, people seem to be attracted to such churches because they feel part of a broader community which looks after them, be it through other church members, church leaders, or God. When faced with any personal or family problems, 71% of our participants would call their pastor and 25% would ask another church member for help. This is consistent with survey evidence from the Afrobarometer and the World Values Survey showing that across Africa, religious leaders are considered amongst the most trustworthy members of civil society. They are expected to take responsibility for their members' welfare in the absence of government-led social interventions and urbanization enhances these expectations (McCauley (2013)).

Pentecostal adherents are expected to support their church financially. Giving is often multi-layered and different ways of transferring money might have different motivations. *Tithing*, the practice of giving away a tenth of all income, usually takes the form of a non-anonymous monthly payment to the church for which church members receive a receipt. This type of giving is akin to a membership fee to the church community. On the other hand, there is giving for specific purposes such as *pledges*, which are occasional non-anonymous donations involving large amounts of money, often for investment in church infrastructure. Finally, there are *spontaneous offerings*, made on a more regular basis, which are generally anonymous and the amounts are given unobserved (though during collections in Sunday services the fact of going forward to give may be very visible to a member's friends and family). Opportunities for this form of giving are frequent and contributions are expected beyond the other regular or earmarked forms of giving. This includes among others *seed offerings* - the practice of giving money in anticipation of a future material benefit, or *thanks-*

givings in gratitude for already materialized benefits (Maxwell (1998); Gifford (2004)). The type of giving we observe in our experiment falls into this category of spontaneous giving.

3.2.3 Giving to God to access insurance

Giving to the church might interact with the use of the church as an insurer in a number of ways. Firstly, individuals might give to the church in expectation that the church as an institution would reward this sign of commitment by disbursing funds in times of need. Secondly, individuals might use their public giving to send signals that they are good community members to other church members, and expect that other church members then contribute to help them in times of need. The costs of religious participation can be seen as screening mechanisms to ensure that members are reliable and to prevent free-riding (Iannaccone (1992)). Additionally, the community structure of the group with repeated interactions reduces monitoring costs (Berman (2000)). These two types of community-based insurance could be considered “material” insurance.

In addition to its role as a social network, the church is also believed by its members to be a setting for encounters with the divine. The church, therefore, has a value as an insurer because it facilitates access to an interventionist God who can prevent negative shocks and favour positive ones. We call this type of insurance “spiritual” insurance. This form of insurance is interesting because it does not rely solely on giving to the church but on “giving to God”, which can be interpreted as “sowing on any fertile ground” - undertaking projects that will be blessed by God (Gifford (2004)). This expansive interpretation of “giving to God” allows us to extend our experimental design to investigate how participants react to giving to recipients which are separate from the church, but which still fulfil the spiritual mandate. This allows us to make some distinctions about the different types of insurance experienced in the church. These mechanisms are formalised in the model in section 3.4, and the doctrinal features are discussed here.

The first and most prominent feature is the forthright relationship between giving to God and material well-being. Pentecostal preachers across Africa speak of a God who does not want His people to be poor or to suffer.⁷ This mandate is often described as a variant of the “Prosperity Gospel”, the set of teachings that say that “Christianity has to do with success, wealth, and status” (Gifford (2004)).⁸ And in contrast to other Christian groups where rewards are generally given out

⁷Enoch Adeboye, head of the Lagos-based international Pentecostal megachurch Redeemed Christian Church of God, teaches, “God is not poor at all by any standard...the closest friends of God [in the Bible] were wealthy people... Poverty is a curse and prosperity is not evil” (Ukah (2005) pg 260). Similarly Gifford (2004) pg 50 cites a pastor proclaiming “God is in the business of addition and multiplications; Satan is in the business of subtraction and division.”

⁸While the term “Prosperity gospel” has typically been used to describe teachings from some Evangelical churches in North America, some scholars of African Pentecostal theologies argue that the “African prosperity gospel” is more nuanced, certain aspects of it “are borrowed from external sources; others are home-grown. Still, others are derived from external sources and adapted to fit local contexts” (Kalu (2008))

in the afterlife, it is usually pointed out that the expected timeframe of miracles is short and immediate.⁹ The immediacy of the miraculous work is reflected in the structure of many church services that are themed around the material successes of its members proclaimed in lively testimonies.¹⁰ Pastors emphasize how behaving in a manner that unlocks God’s blessing helps to avoid the risks that get in the way of achieving these successes. Their preaching makes a particularly strong and explicit link between giving to God and insurance: “The Lord will remember your offering; it is comprehensive assurance” (Gifford (2004) Pg 65).¹¹ And unlocking God’s blessings is invariably achieved by “doing something special that will move God and cause him to bless you more than you intended” i.e. giving money, preferably to the church (Ukah (2005)).

3.3 Experiment

We ran an experiment in Accra, Ghana involving 576 participants in 69 sessions, with a median number of participants per session of 10 and a mean of 9.3. Randomisation was done at the session level. At the start of each session, one participant per group was invited to pick one out of three unmarked envelopes. Participants were told that the envelopes contained a piece of paper that would determine the type of session they would play. The pieces of paper were labelled “Insurance”, which corresponded to our insurance enrolment treatment; “Insurance information”, which corresponded to the control, and “No insurance”, which corresponded to the no insurance information treatment (i.e., groups who were not enrolled in the policy and did not have any sort of discussion about insurance). The insurance was a funeral insurance policy offered by a leading micro-insurer active in the Ghanaian market. The outcomes we measured were allocations out of an endowment (which subjects were free to keep, all or in part) towards a participant’s church and two other non-church recipients.

3.3.1 Enrolment treatment: Funeral insurance policy

Funerals are large and costly events in many sub-Saharan African societies (Berg (2016); Case et al. (2013)). Surviving family members are expected to honour the dead through lavish commemorations. The rising toll of funeral costs has received attention from media, academics and

⁹Some declarations made during a church service include: “Before the year comes to an end, there are people here who will be counting millions of dollars in their accounts...”, “This month, no accident sickness, stagnation will affect you” or “Marriages don’t collapse here, businesses don’t collapse here, nobody fails here” (Gifford (2004) Pg 50 and 58)

¹⁰Gifford cites some testimony themes that include the following: “I am a millionaire today!”, “Jobs started coming in!”, “My own choice car!”, “Now they pay me in dollars!”, “A Baby, Two Cars, and Promotion”.

¹¹Other examples include: “The preacher asked the congregation if they were putting aside for a rainy day, but promised: ‘If you give, you won’t have a rainy day’” (Gifford (2004) Pg 50). “Every part of your destiny is assured by God, insured by God” (Gifford (2004) Pg 58).

political leaders.¹² More recently private financial service providers have begun to offer savings and insurance products designed specifically to meet these costs.

In Ghana, guests and other members of the bereaved's community typically make contributions that help to cover the funeral costs. It is important to note that community support is not only financial - churches also organise the provision of food and moral and logistical support, so any formal insurance product will only be addressing a single aspect of the church contributions. The degree of formalisation of this type of support varies across the churches in our sample. In interviews with church leaders, most confirmed that observed commitment from members was a prerequisite for church involvement in their funerals. Definitions of commitment always included attendance of church events and financial commitment to the church in terms of tithes and offerings.

The funeral policy we offered to participants is a micro-insurance offered in the Ghanaian market for pre-existing groups and needs to be coordinated by a member of the group. In our case, the coordinators were the Pastors of each church branch. The policy covered the life of the participant and a member of his or her immediate family. If either of these parties were to pass away within a year, the policy would pay GHS 1000 (\$265) to the surviving family members. This policy cost GHS 12.5 (\$3USD) for two family members per year.¹³ Individuals in this treatment were enrolled on the spot after completing a demographic survey and before playing the dictator game.

3.3.2 Control group and no insurance information treatment

During pre-tests, it became clear that discussions of death and planning around death would be sensitive topics. Reluctance to contemplate large unpleasant risks has been raised in the literature, particularly in other developing country settings where people are severely limited in the steps they can take to address these risks (Case et al. (2013)). Furthermore, findings in psychological research show that awareness of mortality can modify Christians' religiosity and beliefs in a supernatural entity (see for instance Jong et al. (2012); Norenzayan and Hansen (2006)). As the insurance treatment was designed to isolate the effect of being *enrolled* in insurance, we offered the same information about the insurance policy to the control group, so that the same issues of death would be salient in both settings.

Furthermore, we were interested in isolating the potential effect of risk priming which we ex-

¹²See for instance "The long goodbye: Why funerals are big deals in Ghana" By Paula Newton available at <http://edition.cnn.com/2014/03/11/world/africa/on-the-road-ghana-funerals/index.html> and "Bankruptcy and burials: Can Ghanaians afford such splendid funerals?" The Economist 2007 available at <http://www.economist.com/node/9234475>, and also "Death around the world: Funerals and fantasy coffins in Ghana" available at <https://www.funeralzone.co.uk/blog/death-around-the-world-funerals-in-ghana>

¹³One main goal of the group micro-insurance is to have a simplified scheme and low coordination costs. This implied that the tariff was independent of the gender, age or other characteristics of the insured. The main insurance taker only had to be above the age of 18 and below the age of 75 who could then insure his/her children and other close family members under the age of 75.

pected would modify participants' behaviour. Since discussion about risky events is part of church services as described in the context section, participants' reaction to a risk prime can be linked to their giving behaviour in church. We, therefore, added a second treatment where people were not informed about the content of the insurance policy. Comparisons between this group that did not discuss death and the control group that received insurance information allow us to see how church members react to an increase of perceived risk in the absence of formal insurance.

3.3.3 Recruitment

We recruited 576 study participants from different church branches within one particular denomination. Participants for the main study were recruited from six church branches through announcements made on Sunday mornings during regular church services, and to avoid confounds with normal Sunday offering, all sessions took place during the subsequent work week in a neutral location. A subset of 122 participants was inadvertently recruited during "revival weeks" when their churches were engaged in active fundraising services during the work week. We found interesting results for this subset. They are discussed after presenting the main experimental results, which involve 454 participants.

We were also interested in seeing how our hypothesised mechanisms operate within a secular organisation, so we recruited an additional 242 market sellers. Traders in this market are organized into an association that could provide financial assistance such as credit or insurance to dues-paying members. During the first round of data collection, we realized that the insurance treatment did not operate in this sample as it did within the group of church members. Indeed, the funeral insurance was coordinated by the head of the market association and informal discussions with study participants as well as their answers to survey questions warned us that trust in the insurance coordinator might be low. It, therefore, seemed likely to us that participants would not trust that the insurance would be actually implemented, thus discrediting the insurance treatment. Furthermore, the questionnaire answers informed us that the market association is not a commonly-used risk-sharing structure, and is by no means similar to the church community in that respect. We, therefore, stopped collecting data on market members after the first round of the study and the sample for this group is too small to be able to detect treatment effects. Results from this smaller additional sample are available from the authors on request.

3.3.4 Experimental Setup and Design

It was very important for the credibility of our study that, at the exception of recruitment, all interactions with participants took place off church premises and that participants be assured of anonymity so as to avoid any contamination of the results by perceived pressure from the church authorities. This involved a substantial effort to transport the recruited individuals to a study

location at some distance from the church, as well as the setting up of a proper lab-in the-field with laptops and room dividers between subjects so as to make the assurance of anonymity credible. All participants were compensated for transport to the neutral locations.

Interested participants were assigned to sessions of 8 - 12 people, randomized between treatments and control. A session consisted of an extended survey where each participant was interviewed by an enumerator who spoke the participant’s local language and a set of dictator game decisions.

Participants privately played 10 modified dictator games. Each game asked participants to allocate GHS 11 (a little less than average daily income) between two recipients. The set of recipients consisted of the participant’s church, a secular charity, a national prayer organisation, and the possibility of keeping the money. There were also two ways in which individuals would give to the church: the first being an anonymous donation, the second being a named donation. The pairs of recipients are listed in Table 3.1. The order of the 10 dictator games was randomized by the computer program.¹⁴

A	Self	Church (anonymous)
B	Self	Street children
C	Self	Thanksgiving
D	Self	Church (non-anonymous)
E	Church (anonymous)	Street children
F	Church (anonymous)	Thanksgiving
G	Church (anonymous)	Church (non-anonymous)
H	Street children	Thanksgiving
I	Street children	Church (non-anonymous)
J	Thanksgiving	Church (non-anonymous)

Table 3.1: Pairs of dictator game recipients.

The Street Children’s Fund is a charity that takes care of the educational needs of homeless and vulnerable children. The charity operates in a district of the city that is geographically and culturally distinct from the ones where we recruited participants. Giving to this charity could largely be understood as an altruistic action. The thanksgiving offering is part of an annual inter-faith prayer event. Leaders and members of various faiths join together in prayer for Ghana. Giving towards this event was meant to be interpreted as giving towards a largely spiritual interest. Pre-tests and focus groups during piloting confirmed that study participants would see these two recipients in this manner.

¹⁴Regarding anonymous and non-anonymous donations, one concern is that participants that see named donations before anonymous choices may think about the latter in a different way. We randomize the order of the dictator games precisely to address this type of concern. However, we also checked whether seeing named donations before anonymous ones affect later choices to give to church anonymously and we find that this is not the case.

Participants were paid a flat show-up fee of GHS 20. After all decisions had been made in the dictator games, one game was selected at random individually for each participant, and further payments were made according to the decision taken for that game. This meant that participants had the opportunity to earn up to GHS 31. Average overall earnings from the experiment were GHS 22.50 which were paid privately to participants.

The experiment protocol and the questionnaire can be found in the appendix. This experiment was pre-registered in the AEA registry.¹⁵

3.4 Model

We develop a simple model to formalize the types of behaviour we expect our experiment to capture, and use results from the model to motivate the experimental hypotheses. First, we establish how church members who derive utility from secular as well as spiritual consumption would behave given the risk of an exogenous income shock in a bad state of the world if there is no insurance offered by the church. We then consider two church-based insurance channels. Community-based insurance is modelled as a payment given to a church member in the case of a loss where the size of the payment depends on how much the church member gives to their own church. Spiritual insurance is modelled as a belief that the subjective probability of a loss is reduced by giving to the church and other goods used for religious signalling, such as the spiritual and secular charities.

The mechanisms of our experiment are captured in the model as changes to the size of the loss. The insurance scheme into which we enrolled individuals paid out in the event of their own death or the death of a relative. The loss we model is a financial loss and is not intended to cover all of the aspects of the distress suffered by an individual facing the death of a loved one, let alone the adverse emotional and other consequences to them of their own death. As in any life insurance decision, the individual does not expect to be compensated for their own death but rather intends to offset in part the financial consequences of death to those still left alive. Individuals are capable of anticipating the benefit of that payout and considering it as a (possibly vicarious) benefit in anticipation of which they are willing to pay a premium, even if they do not expect to be alive when the payout occurs.

There are two ways in which our insurance treatments affect the size of the loss. The first is when insurance information is provided to participants. The process necessitates a lengthy discussion about death both of the participating individual and his or her loved ones. This discussion brings home to participants the reality of a risk they had perhaps been inclined not to think about, and thereby makes death more salient. We represent this as an increase in the perceived size of the loss, compared to the situation before they had received this information.

¹⁵AEA RCT Registry ID: AEARCTR-0000558

The second way in which our insurance treatment affects the size of the loss is when individuals are actually enrolled. We represent this as a reduction in the perceived size of the loss.

We show how in the presence of an insurance motive in the church, the main insurance treatment can lead to a decrease in giving, and the provision of only insurance information can lead to an increase in giving.

3.4.1 Setup

We assume that a church member has an income of Y and chooses to give an amount g to the church. The church member enjoys utility $u(Y - g)$ from consuming $Y - g \geq 0$ secular goods, and utility $\theta f(g)$ from contributing $g \geq 0$ to church goods. Thus a church member who gives g to the church enjoys a total utility of $u(Y - g) + \theta f(g)$. The parameter θ reflects the relative weight the individual puts on church activities compared to secular consumption. This weight might differ from one individual to the next (there may be individual heterogeneity). More importantly for our empirical analysis, it might also differ in time (for instance, in revival weeks individuals go to church every day and are focused on spiritual activities). Both utility functions u and f are increasing and concave in their arguments. We, therefore, make the assumption that religious consumption is a normal good.¹⁶

In each period, church members face a probability π of an income loss of size D . Under the assumption that insurance is offered through the church, giving an amount g also has the impact of reducing the size of the loss by an amount $l(g)$, thus the total loss would be $D - l(g)$. Under the assumption that church members believe in spiritual insurance, the probability of a loss is decomposed into a basic probability of loss $\tilde{\pi}$, and a reduction in the probability of loss that can be mitigated by giving money to spiritual goods. Therefore, the total subjective probability of loss is $\pi = \tilde{\pi} - p(g)$.

The following subsections set-up the maximization problems and show how optimal giving varies with the perceived size of the loss. All proofs are in the appendix.

3.4.2 Optimal giving to the church in the absence of any insurance

In this section we assume that church members choose a particular level of giving g to maximize their total expected utility. There is no insurance offered through the church. We label this

¹⁶This assumption contradicts the secularization hypothesis that higher income leads to lower religiosity. However, it is in line with new studies using panel or microdata (Becker and Woessmann (2013) and Buser (2015)). A theoretical rationale for the positive association of income on religious consumption can be found in club good models of religion (Iannaccone (1992), Berman (2000), Berman and Laitin (2008)).

Hypothesis H0. The optimization is:

$$\max_g (1 - \pi)u(Y - g) + \pi u(Y - g - D) + \theta f(g) \quad (3.1)$$

Solving for the first order conditions and taking optimal giving g^* as a function of D allows us to show that giving is decreasing with the size of the loss D :

$$\frac{\partial g^*}{\partial D} < 0 \quad (3.2)$$

This classical result of consumption smoothing comes from the standard concavity assumptions of the utilities derived from secular and religious consumption. When faced with an increase in the potential loss D , church members shift spending from religious consumption to secular consumption to ensure a higher level of secular consumption in case of loss. We call this result the “income effect”.

Therefore, this subsection predicts that if there is no insurance mechanism in the church (neither community-based nor spiritual), the information about the funeral policy in the control group, by increasing the perceived loss D , would lead individuals to decrease church giving g^* .

We summarize this empirical prediction of Hypothesis H0 in the form of:

Empirical prediction H0a: Compared to no information, insurance information decreases giving to the church.

Correspondingly, compared to the control group, the enrolment treatment reduces the perceived risk of loss, thereby decreasing D , and thus would increase giving to the church.

The second empirical prediction of Hypothesis H0 is:

Empirical prediction H0b: Compared to insurance information, insurance enrollment increases giving to the church.

3.4.3 Community insurance: optimal giving when giving reduces the size of a loss

In this section, we assume that the church community provides material insurance such that church giving reduces the size of the loss. We call this Hypothesis H1, and we write $L = D - l(g)$, where l is an increasing and concave function with $l(0) = 0$ and $l'(0) > 1$ (this last condition means that the community insurance in the bad state more than offsets the amount previously donated by the individual). In this case, church giving can be seen as payment of the premium of informal insurance that covers part of the possible loss. As it is offered by the church community, only giving to the church - and not giving to other religious goods - provides access to this type of insurance.

The utility maximization problem is as follows:

$$\max_g (1 - \pi)u(Y - g) + \pi u(Y - g - D + l(g)) + \theta f(g) \quad (3.3)$$

We solve for the first order conditions and express optimal giving g^* as a function of the loss D . We find that:

$$\frac{\partial g^*}{\partial D} > 0 \Leftrightarrow l'(g^*) > 1 \quad (3.4)$$

When there is an insurance motive behind church donations, an increase in the potential loss D triggers two opposite effects: the substitution effect, whereby church members try to mitigate the increase in loss by buying more informal insurance; and the income effect, where church members reduce giving to the church to have more money available for secular consumption smoothing.

Condition (3.4) shows that as long as community-based insurance is effective enough in decreasing the loss, which is the case for low enough g^* , the consumption-smoothing effect from the baseline model (income effect) is outweighed by the increased demand for church insurance (substitution effect). Therefore, the overall effect of an increase in D is an increase in optimal giving.

This subsection predicts that if there exists effective community insurance that reduces the size of a loss in case of a shock, the information about insurance should increase church giving (compared to no insurance information) while the enrolment treatment would decrease church giving (compared to the control group). These effects are the *opposite* of the predicted effects discussed in the previous section when there is no insurance motive for giving to the church. There should be no impact of either treatment on giving to other recipients.

The three empirical predictions of Hypothesis H1 are:

Empirical prediction H1a: Compared to no information, insurance information increases giving to the church.

Empirical prediction H1b: Compared to insurance information, insurance enrollment decreases giving to the church.

Empirical prediction H1c: There is no effect of either insurance information or enrolment on giving to the other recipients (the thanksgiving offering and the street children's fund).

3.4.4 Spiritual insurance: optimal giving when giving reduces the subjective probability of a loss

In this section we assume that there is a spiritual insurance motive such that giving reduces the subjective probability of the loss $\pi = \tilde{\pi} - p(g)$. It is important to stress that giving can be to the church or to any other charitable/spiritual organizations that can be used for religious signalling.

The mechanism here works through God: being a good Christian reduces the subjective probability of a negative shock. We call this Hypothesis H2. Utility can now be written as:

$$\max_g (1 - \tilde{\pi} + p(g))u(Y - g) + (\tilde{\pi} - p(g))u(Y - g - D) + \theta f(g) \quad (3.5)$$

We can show that:

$$\frac{\partial g^*}{\partial D} > 0 \Leftrightarrow p'(g^*)u'(Y - g^* - D) > (\tilde{\pi} - p(g^*))u''(Y - g^* - D) \quad (3.6)$$

This condition is harder to interpret intuitively in the non-parametrized form, and for the second order condition to hold requires that the second derivative of the utility function not increase too fast in the size of the loss, which means that the income effect must not be large. So we defer the discussion to the appendix where we investigate it using a CARA utility function. In short, we find that the optimal giving g^* is increasing in the size of loss D when the spiritual insurance is effective enough. Indeed, when this is the case, individuals prefer to invest in decreasing the subjective probability of loss by increasing their religious giving rather than smoothing consumption (in other words, the substitution effect of the loss dominates the income effect).

Therefore, this subsection predicts that providing participants with insurance information would increase giving to any charitable or spiritual organization (compared to no insurance information) while enrolling them in the insurance treatment would decrease giving to any charitable or spiritual organization (compared to the information treatment).

The three empirical predictions of Hypothesis H2 are:

Empirical prediction H2a: Compared to no information, insurance information increases giving to the church.

Empirical prediction H2b: Compared to insurance information, insurance enrollment decreases giving to the church.

Empirical prediction H2c: The effects on giving to the thanksgiving offering or the street children's fund are similar to the effects on giving to the church.

3.4.5 Experimental hypotheses

We collect here all the empirical predictions of the three different insurance hypotheses so they can be compared and tested against one another.

H0 There is no insurance provided through the church (section 3.4.2).

H0a Compared to no information, insurance information decreases giving to the church.

H0b Compared to insurance information, insurance enrollment increases giving to the church.

H1 *Community* insurance is provided through the church (section 3.4.3).

H1a Compared to no information, insurance information increases giving to the church.

H1b Compared to insurance information, insurance enrollment decreases giving to the church.

H1c There is no effect on giving to the thanksgiving offering or the street children’s fund.

H2 *Spiritual* insurance is provided through God (section 3.4.4).

H2a Compared to no information, insurance information increases giving to the church.

H2b Compared to insurance information, insurance enrollment decreases giving to the church.

H2c The effects on giving to the thanksgiving offering or the street children’s fund are similar to the effects on giving to the church.

Our experimental design allows us to test for an insurance mechanism in giving (test for H0a against H1a/H2a and H0b against H1b/H2b) and then test the two insurance channels of spiritual and community-based insurance against each other (H1c against H2c). Spiritual insurance effects should be similar across recipients. However they might be different in magnitude since, for instance, church members could see church donations as the most effective tool to decrease the probability of a loss.

Our model and experimental design do not allow us to make a clear prediction of the result if we compared the group that received no insurance information to the group that received insurance. This is because, as described in the model setup, the insurance treatment is composed of both the actual insurance effect which we are interested in and a salience effect. The model demonstrates how these two effects work in opposing directions. While we are unable to formulate an ex-ante hypothesis about the overall comparison between receiving insurance and not receiving any information, we discuss the results in subsection 3.5.3.3, where we suggest some interpretations about the relative magnitudes we observe and discuss what they might further tell us about the potential mechanisms of church-based insurance.

3.5 Experimental Results

3.5.1 Descriptive statistics

Our main results include 454 church members recruited during regular service weeks from six different church branches and we have additional socio-demographic information for 404 of those.¹⁷

¹⁷The questionnaire software malfunctioned during the first few sessions, so some survey data were not saved. There were no problems with data from the dictator games.

Table 3.6 in the appendix summarises basic demographic variables of these participants. The final column in this table also includes nationally representative demographic information. Consistent with the recruitment process taking place in churches that are not gender-balanced, we find that our study population had more women than men. Only a fifth of our participants had completed at least a high school education. On average, our participants earned approximately GHS 350, equivalent to \$92 dollars, per month or roughly GHS 12, a little more than \$3, per day. Only 31% held National Health Insurance and only 19% held any other sort of insurance prior to participating in the study. Table 3.8 in the appendix presents a summary of these results.

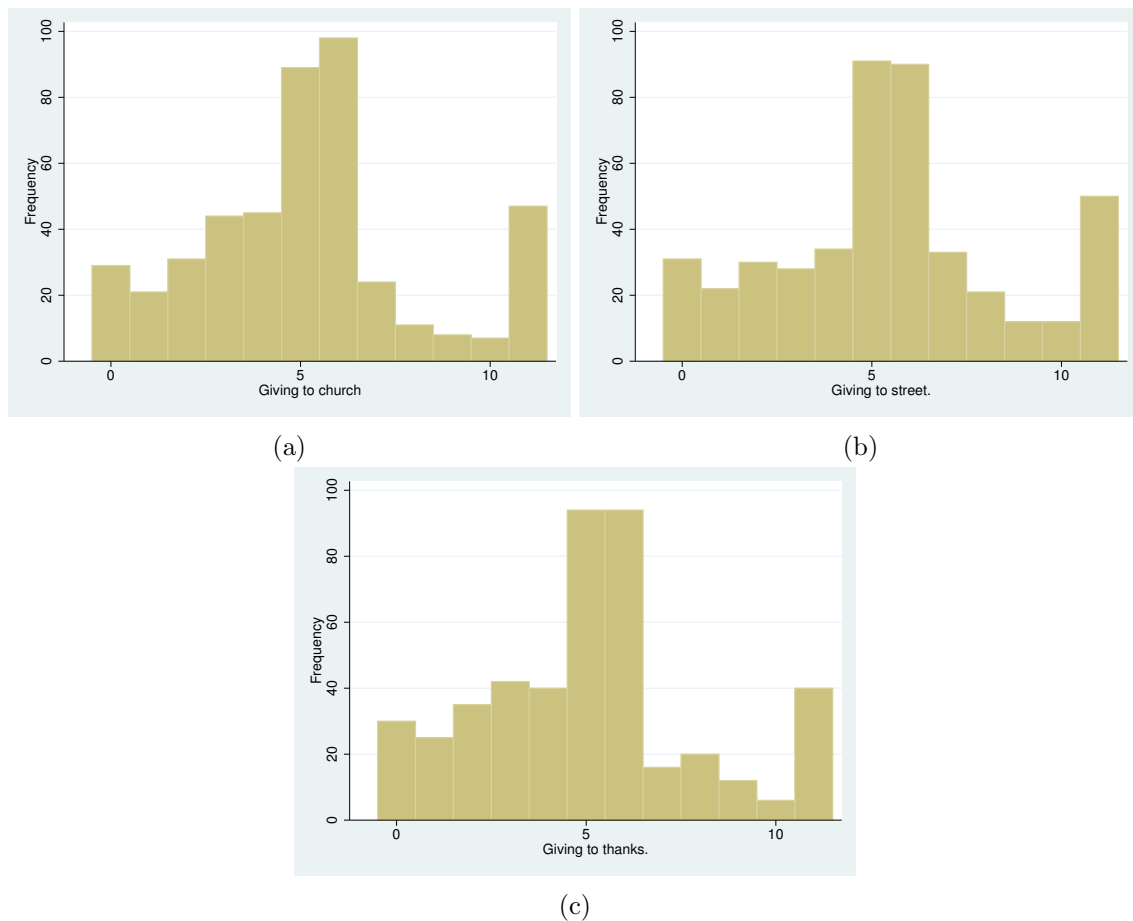
Compared to the national population, our participants had lower incomes and were less likely to be employed. Again, this is consistent with the fact that we recruit relatively active church members. We believe that this selection is not a threat to external validity: we are interested in the attitudes towards insurance in precisely this significant subgroup of the general population who are likely to be particularly vulnerable and face a greater number of formally uninsured risks.

Table 3.7 in the appendix shows that the groups were balanced across treatments for all key variables except gender (women tend to be over-represented in the insurance information and no insurance groups) as shown in Table 3.7. An F-test rejects the hypothesis that these main demographic variables jointly explain the assignment to any of the treatments. We report results of regressions with and without control variables for individuals' demographics (gender, age, education, income, employment status, ethnicity) and measures of religious behaviour (church attendance and prayer frequency) as well as church fixed effects. The treatment effects remain quantitatively similar and their statistical significance is unchanged whether or not we include the demographic variables.

3.5.2 Summary of the allocation decisions

In this subsection and the following, we focus on the three dictator games among the ten played where anonymous donations to the three recipients were paired with the possibility to keep the money. Results from the other games are discussed in subsection 3.5.4. The histograms in Figure 3.1 plot the distributions of giving to the three different recipients. On average, participants chose to keep 5.77 GHS or 52% of their endowment and give 5.22 GHS or 48% of the endowment. We find that 40% of participants selected an allocation of either 5 GHS or 6 GHS. Across recipients, roughly 7% of participants gave nothing, and 10% of participants gave everything away. These spikes at the extreme values highlight that allocations to the recipients may have been censored. To account for this, we report all experimental results using a Tobit regression.

We also find that giving towards the three recipients is significantly correlated, with the correlation coefficients between the pairs of choices ranging from 0.52 to 0.59. If the experimental design induced any order effects, these high degrees of correlation could be problematic for interpretations



Note: Includes participants during normal church weeks (not revival). Shows the amount that participants decided to give to their own church (panel a), the street children's fund (panel b) and the inter-denominational thanksgiving event (panel c) with the alternative of keeping the money.

Figure 3.1: Distribution of giving among normal church population.

across recipients. However, the order of dictator decisions was randomised across participants, mitigating the concern that any order effects could interact with treatment effects.

3.5.3 Treatment effects

3.5.3.1 The effect of insurance information

The first treatment effect tests the hypothesis that a more salient threat of death and a discussion of the associated risk coping strategies affect giving. Table 3.2 presents the basic results of the insurance information on giving relative to the no insurance information treatment (even columns show results of regressions including control variables, odd columns show results of regressions without controls). In the sessions with no insurance information, participants gave an average of GHS 5.05 or 45% of the endowment to the church. Participants who received insurance information increased giving to the church by GHS 0.78, an effect which is statistically significant at 10% and economically large, representing 7% of the endowment and 15% of the mean amount donated by the control groups. This increase in giving is consistent with the hypothesis that the focused discussion of risk puts participants in a fearful state, which makes them more likely to give money to the church. Interestingly there was also an increase in giving of similar magnitude to the street children's fund (an increase of GHS 1.00) and the thanksgiving offering (an increase of GHS 0.75). Tests for equality of the coefficients on the treatment effects across these different outcomes do not reject the null that the difference between them is zero.

	(1)	(2)	(3)	(4)	(5)	(6)
	Giving to church	Giving to church	Giving to thanks.	Giving to thanks.	Giving to street.	Giving to street.
Insurance information	0.778* (0.450)	0.756* (0.412)	0.746* (0.448)	0.662 (0.412)	1.004** (0.457)	1.024** (0.473)
Constant	5.045*** (0.256)	4.561*** (1.403)	4.937*** (0.204)	4.575*** (1.500)	5.220*** (0.266)	5.254*** (1.708)
Individual controls	No	Yes	No	Yes	No	Yes
Church dummies	No	Yes	No	Yes	No	Yes
Observations	289	260	289	260	289	260

Note: Tobit regression censored at 0 and 11. Dependent variables measure donations intended for the own church branch, the street children's fund (secular NGO) and the Inter-denominational Thanksgiving (religious NGO) with the alternative option to keep the money. Control variables include participants demographics and measures of religious behaviour. Standard errors (between parentheses) are clustered at session level. Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 3.2: Giving after receiving insurance information compared to giving with no insurance nor insurance information

The finding that a more salient threat of death increases donations to the church and the other recipients are in line with hypothesis H2a and H2c (Spiritual insurance provided by God).

The GHS 11 endowment used in the dictator games is a little bit more than the median weekly offering to the church. Comparing giving in the no insurance information treatment with the insurance information control, it is interesting to see that a relatively brief discussion about insurance and death could raise giving by 15%, and provides an indication of the importance of the church in this context where there is a lot of uncertainty, but few institutions to deal with it. Furthermore, if this were a pure framing effect unrelated to any insurance motive, we would have expected to see a negative effect on donations to the church instead of the observed positive effect.

3.5.3.2 The effect of being enrolled in an insurance policy

Table 3.3 presents the results of the effect of actually being enrolled in an insurance policy. These results are obtained by comparing people who received insurance with people receiving only insurance information. Column 1 demonstrates that enrolment in the formal insurance policy reduces giving to the church by GHS 0.92. Again, we find similar effects regarding giving to the street children's fund and giving to the thanksgiving offering (a decrease of GHS 0.93 in both cases as shown in column 3 and 5). These numbers correspond to an effect size of approximately 0.55 standard deviations across the three recipients. The effects are once again statistically significant at 10%. They are also economically important in magnitude, representing a reduction in giving equal to over 15% of the mean amount donated in the control group.

Table 3.3: Giving after enrolment in insurance compared to giving with insurance information only

	(1)	(2)	(3)	(4)	(5)	(6)
	Giving to church	Giving to church	Giving to thanks.	Giving to thanks.	Giving to street.	Giving to street.
Insurance	-0.915* (0.471)	-0.698* (0.396)	-0.930* (0.502)	-0.973* (0.497)	-0.926** (0.434)	-1.029** (0.451)
Constant	5.846*** (0.385)	4.733*** (1.501)	5.703*** (0.408)	7.309*** (1.642)	6.227*** (0.377)	7.145*** (1.389)
Individual controls	No	Yes	No	Yes	No	Yes
Church dummies	No	Yes	No	Yes	No	Yes
Observations	285	251	285	251	285	251

Note: Tobit regression censored at 0 and 11. Dependent variables measure donations intended for the own church branch, the street children's fund (secular NGO) and the Inter-denominational Thanksgiving (religious NGO) with the alternative option to keep the money. Control variables include participants demographics and measures of religious behaviour. Standard errors (between parentheses) are clustered at session level. Significance levels: *p<0.1; **p<0.05; ***p<0.01.

These effects are robust when we control for a large set of church and demographic characteristics as demonstrated in columns 2, 4 and 6 of Table 3.3. These controls should pick up variation in church structure on important characteristics such as the level of formal church support during members' funerals and any variation in church teaching on giving. Neither the individual level characteristics nor the church characteristics explain the treatment effects on giving.

Again, these findings are in line with hypothesis H2b and H2c (Spiritual insurance provided by God). Hypothesis H1 - community insurance provided by the church - cannot explain the treatment effects on the non-church recipients.

From this and the previous section, we can see that only Hypothesis H2 - the spiritual insurance hypothesis - is in line with the evidence. Both predictions of Hypothesis H0 - no insurance - contrast with the experimental results, while Hypothesis H1 - community insurance - though compatible with the findings of the treatment effects on giving to the church, is not compatible with the findings on giving to the other recipients.

3.5.3.3 Comparing those being enrolled in the insurance policy with those that do not receive information about the insurance

As mentioned before, the insurance treatment includes both the actual insurance effect and a salience effect and these two effects play in opposite directions. Table 3.9 in the appendix, which compares those being enrolled in the insurance policy with those that do not receive information about the insurance, shows that there is no significant difference in giving levels. Providing insurance information to participants without giving them insurance leads them to increase donations to the church and the two charitable recipients, while the provision of the formal insurance policy leads our participants to decrease their donations to the church and the two charitable recipients. These two effects are of similar magnitude and they cancel out. In other words, providing secular funeral insurance counteracts the effect that a discussion about death has on individuals' behaviour.

The fact that our experiment identifies an enrolment effect that is able to completely offset the salience effect suggests that insurance coverage can provoke immediate behavioural changes. This is in line with recent studies that explore the impact of agricultural insurance policies in developing countries and find strong and immediate responses by farmers regarding production decisions of the insured crop or animal (see Cai et al. (2015), Cole et al. (2017), Karlan et al. (2014)). Provided that they trust the insurance, farmers begin to take riskier decisions such as choosing rain-sensitive crops or a different level of fertilizer shortly after enrollment. The most important behavioural changes in these settings took place on margins associated with the insured products with little change for other uninsured investments, suggesting that risk and insurance coverage (or the lack of it) were salient features in the farmers' decisions.

This does not imply that in ordinary life church members are thinking specifically about insurance against funeral expenses when giving donations to the church. Indeed, in many of the situations when church members make donations, as members of this group do in our dictator games, these are unlikely to be mentally earmarked with the specific goal of decreasing the risks of death. However, the behaviour of our two other groups not only shows that making these risks salient has a real effect on donations - but it also tests for the mechanism by which making the risks salient brings about this behavioural change, namely by inducing church members to seek insurance from the church.

This is how communication within the church operates: pastors' sermons frequently make life risks (e.g. risk of death, or of joblessness, or marital failure, etc.) salient while suggesting donations as a means of coping with these risks. In this way, the church sermons themselves frequently combine the two effects that our experiment seeks to disentangle (salience and insurance). An interesting avenue for research would investigate the ancillary comparison separately: what behavioural responses could we expect from insurance when risks are not made salient?

3.5.4 Discussion

3.5.4.1 How does insurance work?

These experimental results point to an interesting relationship between the type of insurance parishioners might believe they receive from the church and their willingness to engage in costly behaviours to signal membership. Firstly, treatment effects are present across the three recipients. As discussed earlier, neither the street children's fund nor the national thanksgiving offering are linked to the participant's church; and the configuration of our experimental laboratory, with physical partitions between the subjects, rules out any signalling motive. If the type of insurance the participant associates with his/her church membership is purely community based, there should not be a treatment effect on giving to these external recipients.

Thus, the fact that we find similar effects of giving to them, of similar magnitude as effects on giving to the church, indicates that a substantial part of the insurance channel works through beliefs that encourage giving as an act of worship to an interventionist god.

To investigate this spiritual insurance mechanism, we also look at the treatment effects in other dictator games played by participants (as described in table 3.1). First, we investigate the possibility to give to the church where the participant's name would be attached to his donation rather than an anonymous giving. Table 3.13 in the appendix shows that the priming and enrolment treatments have qualitatively the same effect on church giving when donations are not anonymous as when they are. The fact that our subjects do not try to signal their generosity through the use

of nominal donations strengthens the interpretation that charitable behaviour is used as a spiritual mechanism to cope with risk.

We also show in table 3.14 in the appendix that both the insurance information and the insurance enrollment treatment do not modify the participants' decisions to allocate money to their own church against charitable recipients (columns 1 to 4), nor the money allocation between the two NGOs (column 5 and 6). Additionally, figure 3.1 shows that distributions to the different recipients are very similar. It seems that the three different beneficiaries are equally important in the participants' mind for coping with risk.

3.5.4.2 Heterogeneous treatment effects: church members during fund-raising events

Up to this point, we discussed results for church members recruited during normal service weeks. After recruitment, we learnt that two churches had hosted revival weeks during the course of our experiments. Revival weeks are special periods of church activity where members are encouraged to attend church daily. The services consist of prayer, teaching, singing, and exhortation to give money to the church.

Kwabena (2015) describes revival meetings as an essential feature of the contemporary Pentecostal liturgy. In his view, Pentecostal teaching is focused on “scriptures applied in ways that encourage members to invest in financial markets, seize opportunities in education, business, politics and entertainment and wherever able, increase their spheres of influence in the world”. Access to these material benefits is accomplished through religious activities including “massive revival meetings, summits and conferences, all day prayer services and all-night prophetic vigils and mass evangelistic crusades”.

In total, 122 church members participated in the experiment while they were in the middle of a revival week.¹⁸ In terms of demographics, we do not find them to be different from members recruited during regular service weeks. However, participants sampled during revival week are more religiously active (see table 3.10 in the appendix). This could be due to selection or, more likely, to a revival week effect. For this sample, we find important differences in treatment effects. After receiving insurance information, revival week members *decreased* giving to the church and after being enrolled in insurance, they *increased* giving to the church (see table 3.11 in the appendix).

Referring back to the model, these results are consistent with interpreting the revival week as an upwards shift of θ , the relative weight in our subjects' utility function of church activities compared to secular ones. As equation (3.38) in the appendix demonstrates, when equilibrium giving is higher than a given threshold, even in the presence of spiritual insurance, church members respond to an

¹⁸Survey answers were saved for 117 individuals

exogenous shock increasing the size of a loss by a decrease in optimal giving. Intuitively, there is a point at which members have already given so much money to the church, that when faced with the prospect of a negative income shock, they prefer to keep money to smooth secular consumption (i.e., when g^* is large, the income effect dominates the substitution effect).

This explanation of the revival week effect is consistent with the types of activities and benefits members are supposed to derive from revival weeks. Additionally, we find suggestive evidence that people who self report to be habitually high givers respond to treatment in the same manner as people who completed the experiment during a revival week (see Table 3.15 in the appendix). In other words, their reaction to our treatments is the reverse of the reaction of people who are not in revival week, suggesting that for high givers the income effect dominates the substitution effect.

The results highlight that the spiritual state of church members (captured by θ in the model) matters. Intensive religious participation seems to decrease the demand for spiritual insurance and one might derive that overall demand for insurance, formal or informal, is low during important religious events when individuals already “feel protected”. This result would point to a degree of substitution between church participation and church giving; again, this calls for further research.

3.6 Conclusion

We conducted a lab-in-the-field experiment with church members from an established Pentecostal church in Accra, Ghana. We find evidence for religious and charitable giving being part of a church member’s risk-coping strategy. This spiritual insurance channel does not contradict the possibility that other church community-based mechanisms exist in parallel. Indeed, survey responses from church members and leaders emphasize the important role the church plays as a financial contributor. However, our experimental findings add nuance to the literature on religious institutions as coordinating platforms by demonstrating that adherents might care at least as much about spiritual insurance (affecting outcomes through signalling to an interventionist God) as they do about material insurance (accessing transfers of goods and services from other church members).

The treatment effects obtained within the church population depend on three important factors. First, Pentecostal churches stress the involvement of God in terms of blessings in everyday life and teach about God rewarding religious and charitable giving. This particular religious discourse makes members of these churches more prone to see charitable behaviour as a means to decrease the risk of bad events happening and to increase the occurrence of good events. Second, trust in insurance is fundamental, especially in a context where formal institutions are generally weak. In our case, the church was used as a coordinator for the insurance scheme and participants seemed to trust the insurance because it was coordinated by their pastor.¹⁹ Finally, our results obviously

¹⁹By contrast when we tried to run a similar experiment with the market association, the insurance enrollment

depend on the absence (or limited presence) of better institutions to deal with risk. As long as these three conditions are met we expect our results to hold. In particular, we believe that our results would hold in other Pentecostal churches and settings where the development of formal insurance is low. Since the focus on beliefs in religious rituals that influence immediate events are common among a variety of religions and faiths in developing countries, it would be interesting to reproduce the experiment in a different religious setting.

The experiment stressed the importance of religion for economic decisions made by individuals in a setting with weak formal institutions. While individuals might go to religious institutions in those settings because they offer risk-mitigating strategies, we show that formal, private insurance can at least partially substitute spiritual based insurance mechanisms. Since the church was used as a credible coordinator for the insurance scheme, we are inclined to see religious institutions in this context as opportunities to spread formal insurance rather than as an obstacle to its development.

failed to yield any effect because market participants did not trust the head of the market association who was chosen to administrate the insurance scheme.

3.7 Appendix

3.7.1 Tables

	(1)
	mean
Age of the church branch	26.17
Number of church members (approx.)	1035.67
Church members have an education level higher than average	0.33
Church members have income higher than average	0.17
Average number attending Sunday service	610.00
Average amount received on a Sunday	1150.00
The church owns its building	0.83
The church owns other properties	0.17
Number of paid staff	6.67
The church has a welfare fund	1.00
Observations	6

Note: Information about each individual church branch can be found in [Table 3.12](#)

Table 3.4: Summary statistics of church branches

	(1)
	mean
<i>I have found (or I am most likely to find) my spouse:</i>	
In church	0.40
Through friends	0.08
Through family	0.10
Through work	0.15
Through school	0.04
Thanks to other social gathering	0.14
On the Internet	0.00
Other	0.04
<i>Do you do business with people from the same church as you?</i>	
Yes, only	0.09
Yes, I try	0.39
No, not important	0.49
No, I avoid	0.01
<i>Who do you call when you need counselling about personal or family issues?</i>	
No one	0.05
Government or NGO social services	0.00
Friends	0.20
Family	0.39
Pastor	0.71
Church member	0.25
Work superior	0.04
Medical professional	0.03
Other	0.06

Note: Includes all participants. For the last question, multiple answers were possible. The questions were answered in a face-to-face interview with a enumerator.

Table 3.5: Social role of the churches

	(1)	(2)
	Study participants	General population
	mean	mean
female	0.61	0.52
married	0.43	0.39
higher education	0.28	0.15
employed	0.62	0.76
income	362.08	445.50

Note: Includes all study participants (column 1). Figures for general population are from Ghana Living Standard Survey Round 6 (Ghana Statistical Service (2014)). Higher education is defined as having completed at least high school.

Table 3.6: Summary statistics of study participants and comparison with general population

	(1)	(2)	(3)
	Insurance	Insurance information	No insurance
	mean	mean	mean
female	0.54 (0.05)	0.64 (0.77)	0.65
age	36.77 (0.45)	35.68 (0.99)	35.69
married	0.45 (0.49)	0.44 (0.66)	0.41
higher education	0.29 (0.94)	0.26 (0.65)	0.29
employed	0.62 (0.61)	0.59 (0.34)	0.65
monthly income	378.97 (0.40)	361.41 (0.72)	346.67
going to church daily	0.08 (0.88)	0.07 (0.69)	0.08
frequent prayer	0.90 (0.30)	0.93 (0.05)	0.86
any insurance	0.44 (0.83)	0.45 (0.71)	0.42
Observations	144	107	153
F stat		.53	.76
p-value		.85	.65

Note: P-values of a t-test of equality of mean between insurance or insurance information and no insurance groups are reported between parentheses. Higher education is defined as having completed at least high school. Frequent prayer is defined as praying at least once per day.

Table 3.7: Summary statistics of study participants across treatments

	(1)
	mean
<i>Main reason for your church choice:</i>	
Teaching about God	0.87
Moral guidance	0.53
Atmosphere of the services	0.38
Friends and relatives	0.28
Welcoming community	0.21
Presence of interesting and succesful people	0.08
Meeting good marriage partner	0.04
Location of church building	0.15
Comfortable facilities	0.04
Other	0.04
<i>Who would go for financial help?</i>	
No one	0.09
Government or NGO social services	0.02
Friends	0.26
Family	0.53
Church community	0.25
Bank	0.16
Informal financial services	0.08
Work superior	0.04
Others	0.06
<i>Church assistance</i>	
I have received financial assistance from my church in the last 2 years	0.26
<i>Financial decision</i>	
I usually participate to financial decisions in my household	0.62
I usually participate to financial decisions in my household - women only	0.60
<i>Do you hold other type of insurance?</i>	
Yes, NIH	0.31
Yes, other type of insurance	0.19
Yes, NIH or other insurance	0.44

Note: Includes all study participants. In the first 2 questions, multiple answers were possible.

Table 3.8: Financial role of the churches

	(1)	(2)	(3)	(4)	(5)	(6)
	Giving to church	Giving to church	Giving to thanks.	Giving to thanks.	Giving to street.	Giving to street.
Insurance	-0.116 (0.387)	-0.166 (0.290)	-0.159 (0.347)	-0.185 (0.292)	0.0827 (0.377)	0.0198 (0.329)
Constant	5.044*** (0.259)	3.957** (1.657)	4.937*** (0.203)	7.568*** (1.535)	5.219*** (0.267)	5.024*** (1.679)
Individual controls	No	Yes	No	Yes	No	Yes
Church dummies	No	Yes	No	Yes	No	Yes
Observations	334	297	334	297	334	297

Note: Tobit regression censored at 0 and 11. Dependent variables measure donations intended for the own church branch, the street children's fund (secular NGO) and the Inter-denominational Thanksgiving (religious NGO) with the alternative option to keep the money. Control variables include participants demographics and measures of religious behaviour. Standard errors (between parentheses) are clustered at session level. Significance levels: *p<0.1; **p<0.05; ***p<0.01.

Table 3.9: Giving after enrolment in insurance compared to giving with no insurance nor insurance information

	(1)	(2)
	Non revival participants	Revival participants
	mean	mean
female	0.61	0.61
age	36.07	34.96
married	0.43	0.32
higher education	0.28	0.35
employed	0.62	0.56
monthly income	362.08	367.24
going to church daily	0.08	0.11
frequent prayer	0.89	0.96
any insurance	0.44	0.48
Observations	404	117
F stat		1.59
p-value		.12

Note: Revival participants are defined as participants that did the experiment during a revival week. Revival weeks are defined as fundraising weeks with a higher number of services (usually one per day). Higher education is defined as completing at least high school. Frequent prayer is defined as praying at least once per day.

Table 3.10: Comparison of regular and revival week participants

	(1)	(2)
	Giving to church	Giving to church
Insurance	-0.911* (0.475)	2.249*** (0.661)
No insurance	-0.878* (0.460)	2.033** (0.956)
Revival week	-2.247*** (0.619)	
Revival week X Insurance	3.065*** (0.824)	
Revival week X No insurance	2.866*** (1.062)	
Constant	4.726*** (0.956)	2.468 (1.877)
Constant	3.447*** (0.151)	3.585*** (0.261)
Individual controls	Yes	Yes
Observations	521	117

Note: Tobit regression censored at 0 and 11. Dependent variables measure donations intended for the own church branch, the street children's fund (secular NGO) and the Inter-denominational Thanksgiving (religious NGO) with the alternative option to keep the money. Standard errors (between parenthesis) are clustered at session level. Significance levels: *p<0.1; **p<0.05; ***p<0.01.

Table 3.11: Treatment effects on church giving for total sample (column 1) and revival week only (column 2)

	(1)	(2)	(3)	(4)	(5)	(6)
Age of the church branch	27	22	21	22	43	22
Number of church members (approx.)	2500	300	500	500	2114	300
Church members have an education level higher than average	Yes	No	No	No	Yes	No
Church members have income higher than average	No	No	No	No	Yes	No
Average number attending Sunday service	1500	250	200	300	1260	150
Average amount received on a Sunday (GHS)		2000	.	300	.	.
The church owns its building	Yes	Yes	Yes	Yes	Yes	Yes
The church owns other properties	No	No	No	No	Yes	No
Number of paid staff	8	4	6	8	13	1
The church has a welfare fund	Yes	Yes	Yes	Yes	Yes	Yes

Note: Each column represents one church where we recruited participants for the study. Information was gathered by semi-structures questionnaires with a pastor, usually the head pastor.

Table 3.12: Information on church branches

	(1)	(2)
	Giving to Church named (vs keep)	Giving to Church named (vs keep)
model		
Insurance	-1.013** (0.473)	-0.928* (0.483)
No insurance	-0.548 (0.493)	-0.607 (0.495)
Constant	5.607*** (0.413)	4.492*** (1.520)
sigma		
Constant	3.574*** (0.203)	3.528*** (0.214)
Individual controls	No	Yes
Church dummies	No	Yes
Observations	454	404

Note: Tobit regression censored at 0 and 11. Dependent variables measure donations intended for the own church branch, the street children's fund (secular NGO) and the Inter-denominational Thanksgiving (religious NGO) with the alternative option to keep the money. Standard errors (between parenthesis) are clustered at session level. Significance levels: *p<0.1; **p<0.05; ***p<0.01.

Table 3.13: Non-anonymous church giving

	(1) Thanks (vs church)	(2) Thanks (vs church)	(3) Street (vs church)	(4) Street (vs church)	(5) Thanks (vs street)	(6) Thanks (vs street)
Insurance	-0.082 (0.330)	-0.129 (0.298)	0.006 (0.271)	0.024 (0.307)	-0.255 (0.266)	-0.389 (0.262)
No insurance	0.093 (0.335)	-0.114 (0.303)	0.169 (0.258)	0.234 (0.300)	-0.127 (0.254)	-0.388 (0.243)
Constant	4.901*** (0.249)	7.658*** (1.050)	5.702*** (0.205)	6.772*** (0.831)	4.841*** (0.185)	6.769*** (0.963)
Individual controls	No	Yes	No	Yes	No	Yes
Church dummies	No	Yes	No	Yes	No	Yes
Observations	454	404	454	404	454	404

Note: Includes only normal study participants (not revival). Tobit regression censored at 0 and 11. Dependent variables measure donations intended for one institution compared to another institution. Thanks = Inter-denominational thanksgiving fund; Street = Street children's fund; church = own church branch (anonymously). Significance levels: *p<0.1; **p<0.05; ***p<0.01.

Table 3.14: Treatment effects for other dictator games

	(1)	(2)	(3)
	Giving to church	Giving to thanks.	Giving to street.
Insurance	-0.955* (0.540)	-1.197** (0.599)	-1.229** (0.487)
No insurance	-1.198** (0.521)	-1.354*** (0.499)	-1.333*** (0.477)
High givers	-0.750 (0.644)	-1.635** (0.656)	-1.307** (0.633)
High givers X Insurance	0.609 (0.872)	1.173 (0.802)	1.290 (0.825)
High givers X No insurance	1.358* (0.788)	2.373*** (0.812)	0.991 (0.815)
Revival week	-1.893*** (0.620)	-2.436*** (0.491)	-2.376*** (0.443)
Revival week X Insurance	2.800*** (0.939)	2.833*** (0.727)	2.872*** (0.526)
Revival week X No insurance	2.293** (1.013)	2.424** (0.989)	2.191*** (0.744)
employed	0.273 (0.313)	0.478 (0.307)	0.0471 (0.318)
log(income)	0.113 (0.197)	-0.177 (0.248)	0.200 (0.208)
going to church daily	-0.128 (0.569)	-0.925* (0.553)	-0.764 (0.602)
going to church more than once per week	-0.0586 (0.414)	-0.785** (0.357)	-0.697 (0.423)
frequent prayer	0.0404 (0.559)	0.787 (0.548)	-0.278 (0.564)
Constant	5.307*** (1.307)	6.699*** (1.366)	6.293*** (1.298)
Observations	514	514	514

Note: Includes all study participants. Tobit regression censored at 0 and 11. Dependent variables measure donations intended for the own church branch, the street children's fund (secular NGO) and the Inter-denominational Thanksgiving (religious NGO) with the alternative option to keep the money. Standard errors (between parenthesis) are clustered at session level. Significance levels: *p<0.1; **p<0.05; ***p<0.01.

Table 3.15: Comparing revival church members and church members giving relatively high amounts to the church

3.7.2 Model Appendix

3.7.2.1 Setup

We assume that a church member has an income of Y and chooses to give an amount g to the church. The church member enjoys utility $u(\cdot)$ from consuming secular goods, and utility $\theta f(\cdot)$ from consuming church goods. Thus a church member who gives g to the church enjoys a total utility of $u(Y - g) + \theta f(g)$. Both utility functions are concave, thrice differentiable, and increasing in their arguments. In each period church members face a probability π of an income loss of size D .

Under the assumption that insurance is offered through the church community, church giving also has the impact of reducing the size of the loss, thus the total loss would be $D - l(g)$. The function $l(g)$ is assumed to be increasing and concave. Under the assumption that church members believe in spiritual insurance, this probability is decomposed into a basic probability of loss $\tilde{\pi}$, and a portion of the loss that can be mitigated by giving money to spiritual goods. Therefore, the total subjective probability of giving is $\pi = \tilde{\pi} - p(g)$. The function $p(g)$ is assumed to be increasing and concave.

3.7.2.2 Optimal giving to the church in the absence of any insurance

In this section we assume that church members choose a particular level of giving to maximise their total expected utility. There is no insurance offered through the church.

$$\max_g (1 - \pi)u(Y - g) + \pi u(Y - g - D) + \theta f(g) \quad (3.7)$$

This leads to the following first order condition:

$$(\pi - 1)u'(Y - g) - \pi u'(Y - g - D) + \theta f'(g) = 0 \quad (3.8)$$

The second order condition is satisfied:

$$(1 - \pi)u''(Y - g) + \pi u''(Y - g - D) + \theta f''(g) < 0$$

To determine how optimal giving varies with D the size of the loss, we rewrite the FOC in terms of $g^*(D, \theta)$,

$$(\pi - 1)u'(Y - g^*(D, \theta)) - \pi u'(Y - g^*(D, \theta) - D) + \theta f'(g^*(D, \theta)) = 0 \quad (3.9)$$

This equation implicitly defines the optimal giving g^* , which is a function of the expected loss D and θ .

We want to know the impact of experimentally manipulating D on the level of giving of individuals, in other terms the sign of $\frac{\partial g^*(D, \theta)}{\partial D}$.

Let $g_D^*(D, \theta) = \frac{\partial g^*(D, \theta)}{\partial D}$. Taking the derivative of the FOC with respect to D leads to the following equality:

$$g_D^*(D, \theta) * [\pi[u''(Y - g^* - D) - u''(Y - g^*)] + u''(Y - g^*) + \theta f''(g^*)] = -\pi u''(Y - g^* - D) \quad (3.10)$$

The right-hand side of the expression is *positive*. Each individual term of the expression multiplied by g_D^* is *negative*. Therefore g_D^* must also be *negative*.

Let us call $g_\theta^* = \frac{\partial g^*(D, \theta)}{\partial \theta}$. We can also show that g_θ^* is positive: a positive shock on the utility from consuming church goods increases church donations. Taking the derivative of the FOC with respect to θ leads to the following equality:

$$g_\theta^*(D, \theta) * [(1 - \pi)u''(Y - g^*) + \pi u''(Y - g^* - D) + \theta f''(g^*)] = -f'(g^*) \quad (3.11)$$

The right-hand side of the expression is *negative*. Each individual term of the expression multiplied by g_θ^* is *negative*. Therefore g_θ^* is *positive*.

3.7.2.3 Community insurance: optimal giving when giving reduces the size of a loss

In this section we assume that giving to church reduces the size of the loss. $L = D - l(g)$ This assumption illustrates the channel of community insurance.

$$\max_g (1 - \pi)u(Y - g) + \pi u(Y - g - D + l(g)) + \theta f(g) \quad (3.12)$$

This leads to the following first order condition:

$$(\pi - 1)u'(Y - g) + \pi(-1 + l'(g))u'(Y - g - D + l(g)) + \theta f'(g) = 0 \quad (3.13)$$

The second order condition is satisfied:

$$(1 - \pi)u''(Y - g) + \pi l''(g)u'(Y - g - D + l(g)) + \pi(l'(g) - 1)^2 u''(Y - g - D + l(g)) + \theta f''(g) < 0$$

Rewriting the FOC in terms of $g^*(D, \theta)$:

$$\begin{aligned} (\pi - 1)u'(Y - g^*(D, \theta)) + \pi(-1 + l'(g^*(D, \theta)))u'(Y - g^*(D, \theta) - D + l(g^*(D, \theta))) \\ + \theta f'(g^*(D, \theta)) = 0 \end{aligned} \quad (3.14)$$

Taking the derivative of the FOC with respect to D leads to the following equality:

$$g_D^{*'} * [(1 - \pi)u''(Y - g^*) + \pi(l'(g^*) - 1)^2u''(Y - g^* - D + l(g^*)) + \pi l''(g^*)u'(Y - g - D + l(g^*)) + \theta f''(g^*)] = \pi(l'(g^*) - 1)u''(Y - g^* - D + l(g)) \quad (3.15)$$

On the right-hand side $u''(Y - g^* - D + l(g))$ is always *negative* while the expression multiplied by $g_D^{*'}$ of the left-hand side is also always *negative*. Therefore the sign of $g_D^{*'}$ depends on $(l'(g^*) - 1)$.

This provides a relationship between the efficiency of community based insurance and the optimal response of giving.

$$g_D^{*'} > 0 \text{ when } l'(g^*) > 1, \text{ or } g^* < l'^{-1}(1) \quad (3.16)$$

and

$$g_D^{*'} \leq 0 \text{ when } l'(g^*) \leq 1, \text{ or } g^* \geq l'^{-1}(1) \quad (3.17)$$

These conditions tell us that for low levels of optimal giving, experimentally increasing the perceived loss D will decrease optimal giving.

We show below that the variation in optimal giving g^* can be the result of a variation in θ . More particularly, we demonstrate that $g_\theta^{*'} > 0$. Taking the derivative of the FOC with respect to θ gives:

$$g_\theta^{*'} * [(1 - \pi)u''(Y - g^*) + \pi l''(g^*)u'(Y - g^* - D + l(g^*)) + \pi(-1 + l'(g^*))^2u''(Y - g^* - D + l(g^*)) + \theta f''(g^*)] = -f'(g^*) \quad (3.18)$$


The right-hand side of the expression is *negative*. Each individual term of the expression multiplied by $g_\theta^{*'}$ is *negative*. Therefore $g_\theta^{*'}$ is *positive*.

Therefore our model predicts that there exist a threshold level for θ that will trigger a switch in the sign of $g_D^{*'}$.

3.7.2.3.1 Numerical illustration

Graphs 3.2 and 3.3 illustrate a numerical example, in which we simulate the case of $g_D^{*'}(D, \theta)$ changing sign around the threshold $\tilde{\theta}$.

This example uses a CARA utility function ($u(c) = 1 - \exp(-ac)$), and assumes that $l(g) = s \log(1 + g)$. The parameters D, s are chosen such that $D - l(g) \geq 0$. Figure 3.2 shows that the optimal giving $g^*(D, \theta)$ is increasing in θ while Figure 3.3 indicates that $g^*(D, \theta)$ is an increasing



Note: $Y = 10$, $D = 8$, $\pi = 0.4$, $u(\cdot)$ CARA with $a = 0.1$, $f(\cdot)$ CARA with $a = 1$, and $l(\cdot)$ a logarithmic function with $s = 3.5$

Figure 3.2: Community insurance - Numerical example

function of D until the threshold $\tilde{\theta} = 0.345$.

3.7.2.4 Spiritual insurance: optimal giving when giving reduces the subjective probability of a loss

In this section we assume that giving reduces the size of the loss. $\pi = \tilde{\pi} - p(g)$ This assumption illustrates the channel of spiritual insurance.

$$\max_g (1 - \tilde{\pi} + p(g))u(Y - g) + (\tilde{\pi} - p(g))u(Y - g - D) + \theta f(g) \quad (3.19)$$


This leads to the following first order condition:

$$p'(g)u(Y - g) - (1 - \tilde{\pi} + p(g))u'(Y - g) - p'(g)u(Y - g - D) - (\tilde{\pi} - p(g))u'(Y - g - D) + \theta f'(g) = 0 \quad (3.20)$$

Taking the second derivative of the objective function leads to:

$$p''(g)(u(Y - g) - u(Y - g - D)) + 2p'(g)(u'(Y - g - D) - u'(Y - g)) + (1 - \tilde{\pi} + p(g))u''(Y - g) + (\tilde{\pi} - p(g))u''(Y - g - D) + \theta f''(g) \quad (3.21)$$

The first, third, fourth, and fifth terms are negative. The second term is positive. Intuitively, this expression is negative and the SOC is satisfied if the effect of the loss on the marginal utility of income not be too large. In other words, the income effect of the loss on the demand for insurance should not outweigh the substitution effect.



Note: $g^*(D)$ as a function of θ - numerical example marking the threshold $\tilde{\theta}$ where the sign of $g^*(D)$ changes from positive to negative

Figure 3.3: Community insurance - Threshold on θ

This condition is not too demanding. For example, it is satisfied in the numerical simulation shown below, where u is a CARA utility function: $u(c) = 1 - e^{-ac}$. Indeed, in this case,

$$\begin{aligned}
 & p''(g)(u(Y-g) - u(Y-g-D)) + 2p'(g)(u'(Y-g-D) - u'(Y-g)) + \\
 & \quad (1 - \tilde{\pi} + p(g))u''(Y-g) + (\tilde{\pi} - p(g))u''(Y-g-D) + \theta f''(g) \\
 & = [p''(g) + 2ap'(g) - a^2(\tilde{\pi} - p(g))](e^{-a(Y-g-D)} - e^{-a(Y-g)}) - a^2e^{-a(Y-g)} + \theta f''(g) \quad (3.22)
 \end{aligned}$$

A sufficient condition for this expression to be negative is thus given by

$$p''(g) + 2ap'(g) - a^2(\tilde{\pi} - p(g)) \leq 0$$

With $p(g) = k \log(1+g)$ as in the numerical example below, this is equivalent to

$$k \leq \frac{\tilde{\pi}}{\frac{2}{a(1+g)} - \frac{1}{a^2(1+g)^2} + \log(1+g)}$$

This condition holds for example in the simulation below, where $Y = 10$, $D = 8$, $\tilde{\pi} = 0.4$, $a = 0.1$ and $k = 0.09$.

In order to determine how optimal giving varies with D the size of the loss, we rewrite the FOC

in terms of $g^*(D, \theta)$,

$$\begin{aligned} p'(g^*(D, \theta))u(Y - g^*(D, \theta)) - (1 - \tilde{\pi} + p(g^*(D, \theta)))u'(Y - g^*(D, \theta)) - \\ p'(g^*(D, \theta))u(Y - g^*(D, \theta) - D) - (\tilde{\pi} - p(g^*(D, \theta)))u'(Y - g^*(D, \theta) - D) + \theta f'(g^*(D, \theta)) = 0 \end{aligned} \quad (3.23)$$

Taking the derivative of the FOC with respect to D leads to the following equality:

$$\begin{aligned} g_D^{*'} * [p''(g^*)[u(Y - g^*) - u(Y - g^* - D)] + 2p'(g^*)[u'(Y - g^* - D) - u'(Y - g^*)] \\ + (\tilde{\pi} - p(g^*))u''(Y - g^* - D) - u''(Y - g^*)] + u''(Y - g^*) + f''(g) \\ = -[p'(g^*)u'(Y - g^* - D) + (\tilde{\pi} - p(g))u''(Y - g^* - D)] \end{aligned} \quad (3.24)$$

Therefore, we have the following conditions:

$$g_D^{*'} > 0 \text{ when } -[p'(g^*)u'(Y - g^* - D) + (\tilde{\pi} - p(g))u''(Y - g^* - D)] < 0 \quad (3.25)$$

$$g_D^{*'} < 0 \text{ when } -[p'(g^*)u'(Y - g^* - D) + (\tilde{\pi} - p(g))u''(Y - g^* - D)] > 0 \quad (3.26)$$

3.7.2.4.1 Numerical illustration

We simplify these conditions using a CARA utility function: $u(c) = 1 - e^{-ac}$, $u'(c) = ae^{-ac}$, $u''(c) = -a^2e^{-ac}$ and the risk aversion $R(c) = -\frac{u''(c)}{u'(c)} = a$

$$-p'(g) - (\tilde{\pi} - p(g))\frac{u''(Y - g - D)}{u'(Y - g - D)} = -p'(g) + (\tilde{\pi} - p(g))a \quad (3.27)$$

Therefore:

$$g_D^{*'} > 0 \text{ when } -p'(g^*) + (\tilde{\pi} - p(g^*))a < 0 \quad (3.28)$$

$$g_D^{*'} < 0 \text{ when } -p'(g^*) + (\tilde{\pi} - p(g^*))a > 0 \quad (3.29)$$

which can be rewritten as:

$$g_D^{*'} > 0 \text{ when } \frac{1}{a}p'(g^*) > (\tilde{\pi} - p(g^*)) \quad (3.30)$$

$$g_D^{*'} < 0 \text{ when } \frac{1}{a}p'(g^*) < (\tilde{\pi} - p(g^*)) \quad (3.31)$$

Therefore, we find that $g_D^{*'} is positive when the effectiveness of the spiritual insurance divided by the coefficient of absolute risk aversion at g^* is greater than the level of risk at g^* .$

We can now also derive the conditions under which $g_D^{*'} is increasing until a certain level, and then decreasing. For this, we use the following reformulation of conditions (3.30) and (3.31):$

$$g_D^{*'} > 0 \text{ when } p'(g^*) + ap(g^*) > a\tilde{\pi} \quad (3.32)$$

$$g_D^{*'} < 0 \text{ when } p'(g^*) + ap(g^*) < a\tilde{\pi} \quad (3.33)$$

Let us define $\Gamma(g) = p'(g) + ap(g)$. For $g_D^{*}' to be first positive and then negative we need $\Gamma(g)$ to be decreasing:$

$$g_D^{*'} > 0 \text{ when } g^* < \Gamma^{-1}(a\tilde{\pi}) \quad (3.34)$$

$$g_D^{*'} < 0 \text{ when } g^* > \Gamma^{-1}(a\tilde{\pi}) \quad (3.35)$$

In order for Γ' to be decreasing, we need the following condition to be true:

$$\Gamma'(g) \leq 0 \Leftrightarrow p''(g) + ap'(g) \leq 0 \quad (3.36)$$

In the following, we will use a parametrization of $p(g)$ that is concave, and an a such that condition (3.36) hold in order to illustrate that with an increase in the level of giving, due to a higher θ for example, the sign of $g_D^{*}' can reverse.$


We know that g^* is a function of θ , and we will now show a numerical example that illustrates the possibility of $g_D^{*'}|_{g^*(\theta)}$ to be positive until $g^*(\tilde{\theta})$ and negative afterwards. We will use the a simple logarithmic function $p(g) = k \log(g + 1)$ where $\tilde{\pi}$ and k are such that $0 < \tilde{\pi} - p(g) < 1$. If we insert this into equation (3.30), we get:

$$-\frac{k}{g^*(\theta) + 1} + (\tilde{\pi} - k \log(g^*(\theta) + 1))a < 0 \quad (3.37)$$

Together with (3.31), we know that at a specific $\tilde{\theta}$, this equation is equal to zero:


$$-\frac{k}{g^*(\tilde{\theta}) + 1} + (\tilde{\pi} - k \log(g^*(\tilde{\theta}) + 1))a = 0 \quad (3.38)$$

The following graphs illustrate a numerical example, in which we simulate the case of $g^*(\theta)$ being around the threshold in (3.38).



Note: $Y = 10$, $D = 8$, $\pi = 0.4$, $u(\cdot)$ CARA with $a = 0.1$, $f(\cdot)$ CARA with $a = 1$, and $p(\cdot)$ a logarithmic function with $k = 0.09$

Figure 3.4: Spiritual insurance - Numerical example



Note: Marks the threshold $\tilde{\theta}$ where the sign of $g^*(D)$ changes from positive to negative

Figure 3.5: Spiritual insurance: Response of optimal giving to an increase in D

3.7.3 Questionnaire

1. How old are you?
 - number between 18 and 75
2. What is your gender?
 - Male
 - Female
 - prefer not to answer
 - do not know
 - not applicable
3. Where were you born?
 - Accra
 - Rural Ghana
 - Urban Ghana (not Accra)
 - Outside Ghana in Africa
 - Outside Africa
 - prefer not to answer
 - do not know
 - not applicable
4. How long have you lived in Accra?
 - Whole life
 - More than 10 years
 - 5 – 10 years
 - 2-5 years
 - Less than 2 years
 - prefer not to answer
 - do not know
 - not applicable
5. Are you married?
 - Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
6. What is the highest level of education you have completed?
 - No schooling
 - Primary
 - JHSS
 - SHS
 - Polytechnic
 - Professional degree
 - First degree
 - Advanced degree
7. As an adult, have you ever left Ghana for more than 3 months for work or study?
 - prefer not to answer
 - do not know
 - not applicable
 - Yes
 - Where did you go?
 - No
 - prefer not to answer
 - do not know
 - not applicable
8. What is your current work status?
 - Student
 - Unemployed or casual workers
 - Employed
 - Self-employed
 - Inactive(e.g. housewife)
 - Retired
 - prefer not to answer
 - do not know
 - not applicable
9. In which sector is your principal activity?
 - Unemployed
 - Student
 - Government
 - Agriculture
 - Services
 - Manufacturing or construction
 - Not-for-profit
 - prefer not to answer
 - do not know
 - not applicable
10. What is your principal source of income?
 - Salaried job
 - Entrepreneur
 - Family
 - Pension
 - Social support
 - prefer not to answer
 - do not know
 - not applicable
11. How many people live in your household?
 - number

- prefer not to answer
 - do not know
 - not applicable
12. Who usually makes the final financial decisions in your household?
- Me
 - Spouse
 - Parent
 - Other senior relative
 - Joint decisions incl. me
 - Joint decisions not incl. me
 - prefer not to answer
 - do not know
 - not applicable
13. How much do you earn on average per month in Ghana cedis?
- Less than 200
 - 200 - 500
 - 500-1000
 - 1000-1500
 - 1500-2000
 - 2500-3000
 - 3000-8000
 - More than 8000
 - prefer not to answer
 - do not know
 - not applicable
14. How much does your household earn on average per month in Ghana cedis?
- Less than 200
 - 200 - 500
 - 500-1000
 - 1000-1500
 - 1500-2000
 - 2500-3000
 - 3000-8000
 - More than 8000
 - prefer not to answer
 - do not know
 - not applicable
15. Are you responsible for the daily financial needs (not including any salary payments) of any of the following people other than yourself?
- Minor children
 - Adult children
 - Elderly parent
 - Other family
 - Other community member
 - prefer not to answer
 - do not know
 - not applicable
16. How many people in each of the categories are you financially responsible for?
17. Which of the following expenses in you and your dependents' lives do you face on a regular monthly basis?
- Food
 - Sickness(including medication costs, etc.)
 - Transport
 - Insurance
 - Helping needy family members
 - Church contribution
 - Other social contribution (e.g. pledge to a school or charity)
 - Investment in own business
 - Investment in friend, relative, or business partner's business
 - prefer not to answer
 - do not know
 - not applicable
18. How much do you spend each month in cedis on the particular expenses?
19. In the last year, was there a substantial, unexpected increase in any of these aspects of you and your dependant's lives?
- Food
 - Sickness(including medication costs, etc.)
 - Transport
 - Insurance
 - Helping needy family members
 - Church contribution
 - Other social contribution (e.g. pledge to a school or charity)
 - Investment in own business
 - Investment in friend, relative, or business partner's business
 - I faced no unexpected expenses
 - prefer not to answer
 - do not know
 - not applicable

20. How much did you spend to address the substantial increase?
21. What is your principal daily means of transport?
- Privately owned car
 - Privately owned motorcycle
 - Trotro
 - Taxi
 - Bicycle
 - Foot
 - Train
 - prefer not to answer
 - do not know
 - not applicable
22. How many cellphones do you own?
- number
 - prefer not to answer
 - do not know
 - not applicable
23. How many hours a week do you spend browsing the internet?
- Less than 30 minutes
 - 30 minutes to 2 hours
 - 2 to 5 hours
 - 5-10 hours
 - More than 10 hours
 - prefer not to answer
 - do not know
 - not applicable
24. Do you own a business?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
25. How many years have you owned your business?
- Less than 1 year
 - 1 – 3 years
 - 3 – 10 years
 - More than 10 years
 - prefer not to answer
 - do not know
 - not applicable
26. How many employees do you have?
- number
 - prefer not to answer
 - do not know
 - not applicable
27. Which religion do you belong to?
- Pentecostal or charismatic Christian
 - Catholic Christian
 - Traditional
 - Protestant Christian
 - Muslim
 - No religion
 - prefer not to answer
 - do not know
 - not applicable
28. Were you born into this religion?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
29. Which denomination do you belong to? (Please specify a particular denomination, e.g. Action Chapel, ICGC, etc.)
- name
 - prefer not to answer
 - do not know
 - not applicable
30. Do you belong to a different denomination than 5 years ago?
- No
 - Yes, I changed within the last 5 years
 - Yes, I changed more than 5 years ago
 - prefer not to answer
 - do not know
 - not applicable
31. What are the main reasons you are with your current church?
- The teaching about God corresponds to what I believe in
 - I go for the moral guidance to me and my family
 - I like the atmosphere of the services
 - Friends or relatives brought me there

- Other members made an effort to welcome me
 - The congregation contains many interesting and successful people
 - I hope to meet a good marriage partner for me or my children
 - The building is close to my home
 - The facilities are comfortable (e.g. airconditioning, comfortable seating, etc.)
 - prefer not to answer
 - do not know
 - not applicable
32. Are you engaged in any of the following ministries of your church?
- Ushering or welcoming guests
 - Children's ministry
 - Worship team
 - Prayer ministry
 - Men or women's ministry
 - Youth ministry
 - Outreach
 - Deacon or deaconess
 - Protocol
 - Pastoring
 - prefer not to answer
 - do not know
 - not applicable
33. How many hours a week do you spend on each ministry? (preparation, participation, etc.)
34. Have you received financial assistance from your church in the last 5 years (incl. provisions)?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
35. How many times per week or per year do you attend your church?
- Less than once a year
 - A few times per year
 - A few times per month
 - Weekly
- More than once per week
 - prefer not to answer
 - do not know
 - not applicable
36. On average, how many hours do you spend each time you visit your church? (e.g duration of a service, prayer meeting, etc.)
- Less than 1 hour
 - 1-2 hours
 - 2-3 hours
 - More than 3 hours
 - prefer not to answer
 - do not know
 - not applicable
37. How many hours do you travel (going and coming) to attend your regular church?
- Less than 30minutes
 - 30minutes - 1 hour
 - More than 1 hour
 - prefer not to answer
 - do not know
 - not applicable
38. Have you moved your place of residence in order to be closer to your church?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
39. In the last 6 weeks, have you engaged in any of the following activities to bring others to your church?
- Preaching or teaching in the church
 - Preaching or teaching outside the church
 - Inviting friends or family to church
 - Distributing church material in public
 - Speaking or writing in public media (newspapers, radio, television, etc.)
 - Prayer
 - prefer not to answer
 - do not know
 - not applicable
40. How often do you pray to God?

- Multiple times per day
 - Once per day
 - A few times per week
 - Occasionally
 - prefer not to answer
 - do not know
 - not applicable
41. Are there any foods you do not eat for religious reasons?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
42. Do you drink alcohol?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
43. Is this decision for religious reasons?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
44. Do you smoke tobacco?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
45. Is this decision for religious reasons?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
46. Do your religious beliefs affect the way you dress?
- Yes
 - No
 - prefer not to answer
- do not know
 - not applicable
- do not know
 - not applicable
47. Which of the following ceremonies have you done or would you do for your newborn child?
- Traditional naming
 - Church blessing
 - Baptism
 - Outdooring party
 - Nothing
 - prefer not to answer
 - do not know
 - not applicable
48. What was the main cause of the most recent death in your extended family?
- Accident
 - Illness
 - Violence
 - Old age
 - prefer not to answer
 - do not know
 - not applicable
49. Do you think there was spiritual element involved?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
50. How much on average do you give to the church per month?
- Less than 5 cedis
 - 5 – 10 cedis
 - 11 – 50 cedis
 - 51 – 100 cedis
 - 101 – 200cedis
 - 201-500 cedis
 - More than 501 cedis
 - prefer not to answer
 - do not know
 - not applicable
51. Does giving to charity serve the same spiritual duty as giving directly to the church?
- No, charity is more important
 - No, charity is less important

- Yes, they are equally important
 - I do not think I have a duty to either.
 - prefer not to answer
 - do not know
 - not applicable
52. How is God involved in your finances?
- God leaves me to run my own financial affairs
 - God provides enough that I do not suffer
 - God blesses me with financial abundance
 - God is not interested in my finances
 - prefer not to answer
 - do not know
 - not applicable
53. Is it important for you that your close friends come from the same church as you?
- Yes, I try to only make friends with people from my church or mosque
 - Yes, I try to seek people from my church or mosque but it is not so important
 - No, it is not important at all
 - No, I prefer not to be friends with people from my church or mosque
 - prefer not to answer
 - do not know
 - not applicable
54. Is it important for you that your coworkers come from the same church as you?
- Yes, I only work with people from my church or mosque
 - Yes, I try to seek people from my church or mosque but it is not so important
 - No, it is not important at all
 - No, I prefer not to work with people from my church or mosque
 - prefer not to answer
 - do not know
 - not applicable
55. Do you try to do business with people from the same church as you?
- Yes, I only work with people from my church
- Yes, I try to seek people from my church but it is not so important
 - No, it is not important at all
 - No, I prefer not to work with people from my church or mosque
 - prefer not to answer
 - do not know
 - not applicable
56. Who do you call when you need counselling about personal or family issues?
- No one
 - Government or NGO social services
 - Friends
 - Family
 - Pastor
 - Imam
 - Church member
 - Work superior
 - Medical professional
 - prefer not to answer
 - do not know
 - not applicable
57. Who do you go to when you need financial help?
- No one
 - Government or NGO social services
 - Friends
 - Family
 - Church community
 - Bank
 - Informal financial services (e.g susu group, moneylenders)
 - Work superior
 - prefer not to answer
 - do not know
 - not applicable
58. Who do you go to for medical support when you are sick?
- No one
 - Government hospital
 - Private hospital
 - Traditional healer
 - Local clinic
 - Friends
 - Family
 - Pastor

- Imam
 - Work superior
 - prefer not to answer
 - do not know
 - not applicable
59. Within the last 2 years, have you ever attended a prayer camp either for yourself or on behalf of a friend or family member?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
60. Which other clubs or associations are you an active member of?
- name
 - prefer not to answer
 - do not know
 - not applicable
61. How did you meet your spouse if you are married? Where is the most likely place you will meet your spouse if you are not married?
- Church
 - Through friends
 - Through family
 - Through work
 - Through school
 - Other social gathering
 - Internet
 - prefer not to answer
 - do not know
 - not applicable
62. Generally speaking, would you say other Ghanaians can be trusted?
- People can almost always be trusted
 - People can usually be trusted
 - You usually cannot be too careful dealing with people
 - You always cannot be too careful dealing with people
 - prefer not to answer
 - do not know
 - not applicable
63. How much of the time do you trust the Ghanaian government?
- I always trust the government
 - I mostly trust the government
 - I mostly mistrust the government
 - I always mistrust the government
 - prefer not to answer
 - do not know
 - not applicable
64. Are you registered for the National Health Insurance Scheme?
- No
 - Yes
 - prefer not to answer
 - do not know
 - not applicable
65. Why not?
66. Do you hold any other sorts of insurance?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
67. Please specify which ones
- name
 - prefer not to answer
 - do not know
 - not applicable
- ADDITIONAL QUESTIONS FOR MARKET MEMBERS**
68. Do you belong to the Kaneshie Markets Association?
- Yes
 - No
 - prefer not to answer
 - do not know
 - not applicable
69. How many years have you belonged to the association?
- number
 - prefer not to answer
 - do not know
 - not applicable

70. How many years have you paid your market dues?
- number
 - prefer not to answer
 - do not know
 - not applicable
71. Have you ever met Auntie Eva, the market secretary?
- Yes
 - No, but I know who she is
 - No, I have no idea who she is
 - prefer not to answer
 - do not know
 - not applicable
72. Have you ever received financial help from the market (loan, emergency money, help with stock, etc.)?
- Yes, from the head of the market line
 - Yes, from the Kaneshie Market Association
 - Yes, from any other market members
 - No, I have received no help from the market
 - prefer not to answer
 - do not know
 - not applicable
73. Which market line do you belong to?
- name
 - prefer not to answer
 - do not know
 - not applicable

3.7.4 Experimental protocol

Before participants arrive, enumerators prepare the room: they install chairs, tables, computers, and dividers. The session monitor prepares desks numbers to be given to participants, receipts, insurance forms, and envelopes for random draw. Enumerators sit and start the program while waiting for the participants.

When participants arrive, they are given an ID number and sit in front of a computer. The session monitor welcomes them and explains the experiment in English. One enumerator translates the explanations in the local language spoken by the participants (most often Twi).

Session monitor: *“Welcome to our study, thank you for coming. We are a group of researcher from Toulouse in France and we are doing a study here in Ghana in collaboration with Central University College. The first part of the study consists in a questionnaire that contains demographic questions as well as questions about your religious practices. For completing the survey, you will receive 20GHS. The questionnaire is anonymous, no one will be able to associate you with the questionnaire. We will not record any private information with your answers. You are allowed to skip any question you do not want to answer. If you choose to stop the interview we will give you 5GHS for your effort to come here. In a second part, you will play a game on how to allocate 11GHS. You will be asked 10 questions; in each question you can choose you to allocate money between two different recipients. The list of the different recipients is written on the white board. First, in some occasions, you will have the opportunity to keep money for yourself. Another option will be to give money to the Street’s Children Fund. This is an NGO active in Jamestown that helps children from the streets to access education. You will also be allowed to give to the National Thanksgiving Association, for a National Week of Prayer gathering as many Ghanaians as possible, from different religions, to pray for Ghana. This year the Week of prayer is expected to be held in March. You will also have the chance to make some donation to your own church privately, and in another option you can offer money to your church with your name. At the end of the game, we will pick one question randomly with equal probabilities. We will make payments according to your answers to this question. First, you will play a practice round, helped by your enumerator. You will play the real round privately. The session monitor is here to help you at any time. There are different types of sessions. The session we run with you is picked at random. We now ask you if you agree with us using the data from this session - the answers you give in the questionnaire and in the game which are anonymous - for research purpose.”*

PAUSE, wait for objections. Continue if there are no objections. One participant draws an envelope that contains the type of the session: ‘INSURANCE’, ‘INSURANCE INFORMATION’, or ‘NO INSURANCE’.

If the session is ‘INSURANCE’:

Session monitor: *“The type of session we run with you is called ‘INSURANCE’. This means that we will offer you a funeral insurance policy. This insurance scheme is offered on the market by Enterprise Life and we will purchase it for you. From the day we enroll your name and for one year, if you or one of your family members (to be designated) dies, a recipient will get money to finance the funeral expenses (1000GHS). You can choose to enroll with you one family member. It can be of your parents, one of your minor children, your wife or husband. They should be younger than 74. To cover you and a family member, this insurance costs annually 12.50GHS. This group insurance is coordinated by your church. The coordinator is your Pastor (name). At the end of the year, you can collectively decide to continue or not the insurance contract. Any claim should be addressed to your Pastor, he will receive money for you if anything happens. If you do not wish to be benefit from the insurance policy, please inform us. “*

If the session is ‘INSURANCE INFORMATION’:

Monitor: *“The type of session we run with you is called ‘INSURANCE’. This means that we will give you information about an insurance that you can decide to purchase for yourself. The insurance we would like to discuss with you is a funeral insurance offered by Enterprise Life: if you or one of your family members (to be designated) dies, a recipient will get money to finance the funeral expenses (1000GHS). You can choose to enroll with you one of your parents, one of your minor children, or your wife or husband. They should be younger than 74. To cover you and one family member, this insurance costs annually 12.50GHS. At the end of the year, you can collectively decide to continue or not the insurance contract. The coordinator for this group insurance could be your Pastor (name). Any claim would be addressed to him, he would receive money for you if anything happens. If you are interested in buying this insurance, please contact Pastor (name).”*

For all the sessions:

Monitor: *“The experiment is now about to start. Please remember that the answers that you will give will be anonymous and be kept secret. The enumerators are here to guide you through the survey and translate it if necessary. You can leave the experiment at any time and you will be paid a show-up fee of 5GHS.”*

The participants answer the questionnaire with the help of enumerators. After answering the questionnaire, participants in ‘INSURANCE’ sessions fill in the insurance form. Then they proceed to the game, first a practice round helped by the enumerators, and then the real round by themselves. Once they have entered their choices for the 10 money allocation questions, the participants randomly pick a letter that corresponds to one question. The session monitor enters the letter in the computer to record total payment. Participants leave the room and wait to be called for payment.

Bibliography

- Adachi, H. (2003). A search model of two-sided matching under nontransferable utility. *Journal of Economic Theory*, 113(2):182–198.
- Adams, A. and Andrew, A. (2019). Preferences and beliefs in the marriage market for young brides. Technical report, CEPR Discussion Paper.
- Ager, P. and Ciccone, A. (2016). Agricultural risk and the spread of religious communities. *Working Paper*.
- Almond, D., Li, H., and Zhang, S. (2019). Land reform and sex selection in china. *Journal of Political Economy*, 127(2):000–000.
- Ananat, E. O. and Hungerman, D. M. (2012). The power of the pill for the next generation: oral contraception’s effects on fertility, abortion, and maternal and child characteristics. *Review of Economics and Statistics*, 94(1):37–51.
- Angrist, J. D. and Evans, W. N. (1996). Children and their parents’ labor supply: Evidence from exogenous variation in family size. Technical report, National bureau of economic research.
- Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434):444–455.
- Assemblies of God World Missions (2017). Vital statistics.
- Auriol, E., Lassebie, J., Panin, A., Raiber, E., Seabright, P., et al. (2017). God insures those who pay? formal insurance and religious offerings in ghana. Technical report, Toulouse School of Economics (TSE) Working Paper.
- Banerjee, A., Dufló, E., Ghatak, M., and Lafortune, J. (2013). Marry for what? caste and mate selection in modern india. *American Economic Journal: Microeconomics*, 5(2):33–72.
- Baochang, G., Feng, W., Zhigang, G., and Erli, Z. (2007). China’s local and national fertility policies at the end of the twentieth century. *Population and Development Review*, 33(1):129–148.
- Becker, G. S. (1960). An economic analysis of fertility. In *Demographic and economic change in developed countries*, pages 209–240. Columbia University Press.
- Becker, G. S. (1973). A theory of marriage: Part i. *Journal of Political economy*, 81(4):813–846.
- Becker, G. S. (1981). A Treatise on the Family. *Cambridge/MA*.
- Becker, G. S. and Lewis, H. G. (1973). On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81(2, Part 2):S279–S288.

- Becker, G. S., Murphy, K. M., and Tamura, R. (1990). Human capital, fertility, and economic growth. *Journal of political economy*, 98(5, Part 2):S12–S37.
- Becker, S. O. and Woessmann, L. (2013). Not the opium of the people: Income and secularization in a panel of prussian counties. *American Economic Review*, 103(3):539–44.
- Belot, M. and Francesconi, M. (2006). Can anyone be ‘the’one? evidence on mate selection from speed dating. *Working Paper*.
- Berg, E. (2016). Securing the afterlife: The puzzle of funeral insurance. *Working Paper*.
- Berman, E. (2000). Sect, Subsidy and Sacrifice: an Economist’s View of Ultra-Orthodox Jews. *Quarterly Journal of Economics*, 115(3).
- Berman, E. and Laitin, D. D. (2008). Religion, terrorism and public goods: Testing the club model. *Journal of Public Economics*, 92(10-11):1942–1967.
- Botha, E. (2007). The new reformation: The amazing rise of the pentecostal-charismatic movement in the 20th century. *Church History Society of Southern Africa*.
- Bovet, J., Raiber, E., Ren, W., Wang, C., and Seabright, P. (2018). Parent–offspring conflict over mate choice: An experimental study in china. *British Journal of Psychology*, 109(4):674–693.
- Brandt, L., Li, H., Turner, L., and Zou, J. (2018). Are china’s “leftover women” really leftover? an investigation of marriage market penalties in modern-day china. Technical report, RCEA Working Paper.
- Buser, T. (2015). The effect of income on religiousness. *American Economic Journal: Applied Economics*, 7(3):178–95.
- Cai, H., Chen, Y., Fang, H., and Zhou, L.-A. (2015). The effect of microinsurance on economic activities: evidence from a randomized field experiment. *Review of Economics and Statistics*, 97(2):287–300.
- Case, A., Garrib, A., Menendez, A., and Olgiati, A. (2013). Paying the Piper: The High Cost of Funerals in South Africa. *Economic Development and Cultural Change*, 62(1):1–20.
- Chang, L., Wang, Y., Shackelford, T. K., and Buss, D. M. (2011). Chinese mate preferences: Cultural evolution and continuity across a quarter of a century. *Personality and Individual Differences*, 50(5):678–683.
- Chen, D. (2010). Club Goods and Group Identity: Evidence from Islamic Resurgence during the Indonesian Financial Crisis. *Journal of Political Economy*, 118(2):300–354.
- Chiappori, P.-A. and Salanié, B. (2016). The econometrics of matching models. *Journal of Economic Literature*, 54(3):832–61.
- Chiappori, P.-A., Salanié, B., and Weiss, Y. (2015). Partner choice and the marital college premium: Analyzing marital patterns over several decades. *Working Paper*.
- Choo, E. and Seitz, S. (2013). The collective marriage matching model: Identification, estimation, and testing. *Structural Econometric Models*, pages 291–336.

- Choo, E. and Siow, A. (2006). Who marries whom and why. *Journal of political Economy*, 114(1):175–201.
- Cole, B. S., Giné, X., Tobacman, J., Topalova, P., Townsend, R., and Vickery, J. (2013). Barriers to Household Risk Management: Evidence from Indi. *American Economic Journal: Applied Economics*, 5(1):104–135.
- Cole, S., Giné, X., and Vickery, J. (2017). How does risk management influence production decisions? evidence from a field experiment. *The Review of Financial Studies*, 30(6):1935–1970.
- Cox, H. (2011). Foreward. In Brown, C. G., editor, *Global Pentecostal and Charismatic Healing*, pages xvii–xxii. Oxford University Press.
- Cox, W. (2018). Demographia world urban areas. *14th Annual Edition*.
- Dehejia, R., DeLeire, T., and Luttmer, E. F. (2007). Insuring consumption and happiness through religious organizations. *Journal of Public Economics*, 91(1-2):259–279.
- Del Boca, D. and Flinn, C. J. (2014). Household behavior and the marriage market. *Journal of Economic Theory*, 150:515–550.
- Du, J., Wang, Y., and Zhang, Y. (2015). Sex imbalance, marital matching and intra-household bargaining: Evidence from china. *China Economic Review*, 35:197–218.
- Duflo, E., Dupas, P., and Kremer, M. (2015). Education, HIV, and early fertility: Experimental evidence from Kenya. *The American Economic Review*, 105(9):2757–2797.
- Duflo, E., Dupas, P., and Kremer, M. (2017). The Impact of Free Secondary Education: Experimental Evidence from Ghana. *MIT Working Paper*.
- Ebenstein, A. (2010). The “missing girls” of China and the unintended consequences of the one child policy. *Journal of Human Resources*, 45(1):87–115.
- Engel, J. W. (1984). Marriage in the people’s republic of china: Analysis of a new law. *Journal of Marriage and the Family*, pages 955–961.
- Fan, C. C. and Huang, Y. (1998). Waves of rural brides: Female marriage migration in china. *Annals of the Association of American Geographers*, 88(2):227–251.
- Fisman, R., Iyengar, S. S., Kamenica, E., and Simonson, I. (2006). Gender differences in mate selection: Evidence from a speed dating experiment. *The Quarterly Journal of Economics*, 121(2):673–697.
- Fisman, R., Iyengar, S. S., Kamenica, E., and Simonson, I. (2008). Racial preferences in dating. *The Review of Economic Studies*, 75(1):117–132.
- Fort, M., Schneeweis, N., and Winter-Ebmer, R. (2016). Is education always reducing fertility? evidence from compulsory schooling reforms. *The Economic Journal*, 126(595):1823–1855.
- Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.
- Galichon, A. and Salanié, B. (2015). Cupid’s invisible hand: Social surplus and identification in matching models. *Working Paper available at SSRN 1804623*.

- Gershman, B. (2016). Witchcraft beliefs and the erosion of social capital: Evidence from sub-saharan africa and beyond. *Journal of Development Economics*, 120:182–208.
- Ghana Statistical Service (2014). Ghana living standards survey round 6 (glss 6) main report. Technical report.
- Giesbert, L. and Steiner, S. (2015). Client perceptions of the value of microinsurance: evidence from southern Ghana. *Journal of International Development*, 27(1):15–35.
- Giesbert, L., Steiner, S., and Bendig, M. (2011). Participation in Micro Life Insurance and the Use of Other Financial Services in Ghana. *Journal of Risk and Insurance*, 78(1):7–35.
- Gifford, P. (2004). *Ghana's new Christianity: Pentecostalism in a globalizing African economy*. Indiana University Press.
- Goldin, C. and Katz, L. F. (2002). The power of the pill: Oral contraceptives and women's career and marriage decisions. *Journal of political Economy*, 110(4):730–770.
- Gruber, J. and Hungerman, D. M. (2007). Faith-based charity and crowd-out during the great depression. *Journal of Public Economics*, 91(5-6):1043–1069.
- Guo, R., Li, H., Yi, J., and Zhang, J. (2017). Fertility, household structure, and parental labor supply: Evidence from china. *Journal of Comparative Economics*.
- Hadnes, M. and Schumacher, H. (2012). The gods are watching: An experimental study of religion and traditional belief in burkina faso. *Journal for the Scientific Study of Religion*, 51(4):689–704.
- Han, L., Li, T., and Zhao, Y. (2015). How Status Inheritance Rules Affect Marital Sorting: Theory and Evidence from Urban China. *The Economic Journal*, 125(589):1850–1887.
- He, X. and Zhu, R. (2016). Fertility and female labour force participation: Causal evidence from urban china. *The Manchester School*, 84(5):664–674.
- Higgins, L. T. and Sun, C. (2007). Gender, social background and sexual attitudes among chinese students. *Culture, Health & Sexuality*, 9(1):31–42.
- Higgins, L. T., Zheng, M., Liu, Y., and Sun, C. H. (2002). Attitudes to marriage and sexual behaviors: A survey of gender and culture differences in china and united kingdom. *Sex Roles*, 46(3):75–89.
- Hitsch, G. J., Hortaçsu, A., and Ariely, D. (2010). Matching and sorting in online dating. *The American Economic Review*, 100(1):130–163.
- HSBC (2017). Higher and higher. *The Value of Education*.
- Hu, Y. (2016). Marriage of matching doors: Marital sorting on parental background in china. *Demographic Research*, 35:557–580.
- Huang, F., Jin, G. Z., and Xu, L. C. (2016a). Love, money, and parental goods: Does parental matchmaking matter? *Journal of Comparative Economics*.
- Huang, W., Lei, X., and Sun, A. (2016b). When Fewer Means More: Impact of One-Child Policy on Education of Girls. Technical report, Working paper, Harvard University.

- Huang, W. and Zhou, Y. (2015). One-Child Policy, Marriage Distortion, and Welfare Loss. *IZA Discussion Paper*.
- Hull, P. (2018). Isolateing: Identifying counterfactual-specific treatment effects with cross-stratum comparisons. *Available at SSRN 2705108*.
- Iannaccone, L. R. (1992). Sacrifice and stigma: reducing free-riding in cults, communes, and other collectives. *Journal of Political Economy*, pages 271–291.
- Jacobsen, J., Pearce III, J., and Rosenbloom, J. (1999). The economic consequences of unwed motherhood: Using twin births as a natural experiment. *Journal of Human Resources*, 34(3):449–474.
- Jenkins, P. (2011). *The next Christendom: The coming of global Christianity*. Oxford University Press.
- Jia, R. and Persson, T. (2017). Individual vs. Social Motives in Identity Choice: Theory and Evidence from China. Technical report, 21st Century China Center Research Paper No. 2017-065.
- Jong, J., Halberstadt, J., and Bluemke, M. (2012). Foxhole atheism, revisited: The effects of mortality salience on explicit and implicit religious belief. *Journal of Experimental Social Psychology*, 48(5):983–989.
- Kalemli-Ozcan, S. (2003). A stochastic model of mortality, fertility, and human capital investment. *Journal of Development Economics*, 70(1):103–118.
- Kalu, O. (2008). *African Pentecostalism: an introduction*. Oxford University Press.
- Karlan, D., Osei, R., Osei-Akoto, I., and Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics*, 129(2).
- Kurzban, R. and Weeden, J. (2007). Do advertised preferences predict the behavior of speed daters? *Personal Relationships*, 14(4):623–632.
- Kwabena, A.-G. J. (2015). Pentecostalism and the african christian landscape. In Lindhardt, M., editor, *Pentecostalism in Africa: Presence and impact of pneumatic Christianity in postcolonial societies*, chapter 4, pages 100–114. Brill.
- Lam, D. and Duryea, S. (1999). Effects of schooling on fertility, labor supply, and investments in children, with evidence from Brazil. *Journal of Human Resources*, pages 160–192.
- Lavy, V. and Zablotsky, A. (2015). Women’s schooling and fertility under low female labor force participation: Evidence from mobility restrictions in israel. *Journal of Public Economics*, 124:105–121.
- Li, B. and Zhang, H. (2016). Does Population Control Lead to Better Child Quality? Evidence from China’s One-Child Policy Enforcement. *Journal of Comparative Economics*.
- Logan, J. A., Hoff, P. D., and Newton, M. A. (2008). Two-sided estimation of mate preferences for similarities in age, education, and religion. *Journal of the American Statistical Association*, 103(482):559–569.

- Lundborg, P., Plug, E., Rasmussen, A. W., et al. (2017). Can Women Have Children and a Career? IV Evidence from IVF Treatments. *American Economic Review*, 107(6):1611–1637.
- Manser, M. and Brown, M. (1980). Marriage and household decision-making: A bargaining analysis. *International Economic Review*, pages 31–44.
- Maxwell, D. (1998). Delivered from the spirit of poverty? *Journal of Religion in Africa*, 27.
- McCauley, J. F. (2013). Africa’s new big man rule? Pentecostalism and patronage in Ghana. *African Affairs*, 112(446):1–21.
- McCleary, R. M. and Barro, R. J. (2006a). Religion and economy. *The Journal of Economic Perspectives*, 20(2):49–72.
- McCleary, R. M. and Barro, R. J. (2006b). Religion and political economy in an international panel. *Journal for the Scientific Study of Religion*, 45(2):149–175.
- McCrary, J. and Royer, H. (2011). The effect of female education on fertility and infant health: Evidence from school entry policies using exact date of birth. *American Economic Review*, 101(1):158–95.
- Miller, G. (2010). Contraception as development? New evidence from family planning in Colombia. *The Economic Journal*, 120(545):709–736.
- Mobarak, A. M. and Rosenzweig, M. R. (2013). Informal Risk Sharing, Index Insurance, and Risk Taking. *American Economic Review: Papers and Proceedings*, 103(3):375–380.
- Monstad, K., Propper, C., and Salvanes, K. G. (2008). Education and fertility: Evidence from a natural experiment. *Scandinavian Journal of Economics*, 110(4):827–852.
- Mu, Z. and Xie, Y. (2014). Marital age homogamy in China: A reversal of trend in the reform era? *Social Science Research*, 44:141–157.
- Norenzayan, A. and Hansen, I. G. (2006). Belief in supernatural agents in the face of death. *Personality and Social Psychology Bulletin*, 32(2):174–187.
- Nunn, N. and de la Sierra, R. S. (2017). Why being wrong can be right: Magical warfare technologies and the persistence of false beliefs. Technical report, National Bureau of Economic Research.
- Okyerefo, M. P. K. (2011). Pentecostalism in the city of Accra: a blossom on functional appeal and urban fecundity. *American Journal of Sociological Research*, 1(1):27–34.
- Ong, D. and Wang, J. (2015). Income attraction: An online dating field experiment. *Journal of Economic Behavior & Organization*, 111:13–22.
- Osili, U. O. and Long, B. T. (2008). Does female schooling reduce fertility? Evidence from Nigeria. *Journal of Development Economics*, 87(1):57–75.
- Pew Research Center (2006). Spirit and Power: A 10-country Survey of Pentecostals.
- Pimentel, E. E. (2000). Just how do I love thee?: Marital relations in urban China. *Journal of Marriage and Family*, 62(1):32–47.
- Qian, N. (2009). Quantity-quality and the one child policy: The only-child disadvantage in school enrollment in rural China. Technical report, National Bureau of Economic Research.

- Qin, X., Zhuang, C. C., and Yang, R. (2016). Does the one-child policy improve children’s human capital in urban China? A regression discontinuity design. *Journal of Comparative Economics*.
- Raiber, E. (2017). Expected fertility and educational investment: Evidence from the one-child-policy in china. Technical report, Toulouse School of Economics (TSE).
- Riley, N. E. (1994). Interwoven lives: Parents, marriage, and guanxi in china. *Journal of Marriage and the Family*, pages 791–803.
- Rosenzweig, M. R. (1990). Population growth and human capital investments: theory and evidence. *Journal of Political Economy*, 98(5, Part 2):S38–S70.
- Rosenzweig, M. R. and Schultz, T. P. (1989). Schooling, information and nonmarket productivity: contraceptive use and its effectiveness. *International Economic Review*, pages 457–477.
- Rosenzweig, M. R. and Zhang, J. (2009). Do population control policies induce more human capital investment? Twins, birth weight and China’s “one-child” policy. *The Review of Economic Studies*, 76(3):1149–1174.
- Roth, A. E. and Sotomayor, M. (1992). Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541.
- Scharping, T. (2013). *Birth Control in China 1949-2000: Population policy and demographic development*. Routledge.
- Scheve, K. and Stasavage, D. (2006). Religion and Preferences for Social Insurance. *Quarterly Journal of Political Science*, 1(May):255–286.
- Schultz, T. P. (1997). Demand for children in low income countries. *Handbook of Population and Family Economics*, 1:349–430.
- Tang, C., Zhao, L., and Zhao, Z. (2016). Child labor in china. *China Economic Review*.
- Thelen, F. M. (2017). The explosive growth of pentecostal-charismatic christianity in the global south, and its implications for catholic evangelization. *Homiletic and Pastoral Review*.
- Ukah, A. F. (2005). Those who trade with god never lose: The economics of pentecostal activism in nigeria. *Christianity and social change in Africa: Essays in honor of JDY Peel*, pages 253–274.
- Weber, M. (1905). *The Protestant Ethic and the Spirit of Capitalism*.
- Whyte, M. K., Feng, W., and Cai, Y. (2015). Challenging myths about china’s one-child policy. *The China Journal*, (74):144–159.
- Willis, R. J. (1973). A new approach to the economic theory of fertility behavior. *Journal of Political Economy*, 81(2, Part 2):S14–S64.
- Wong, L. Y. (2003). Structural estimation of marriage models. *Journal of Labor Economics*, 21(3):699–727.
- Xia, P., Tu, K., Ribeiro, B., Jiang, H., Wang, X., Chen, C., Liu, B., and Towsley, D. (2014). Characterization of user online dating behavior and preference on a large online dating site. In *Social Network Analysis-Community Detection and Evolution*, pages 193–217. Springer.

Xia, Y. R. and Zhou, Z. G. (2003). The transition of courtship, mate selection, and marriage in china. *Mate selection across cultures*, pages 231–246.

Zwane, A. P., Zinman, J., Van Dusen, E., Pariente, W., Null, C., Miguel, E., Kremer, M., Karlan, D. S., Hornbeck, R., Giné, X., et al. (2011). Being surveyed can change later behavior and related parameter estimates. *Proceedings of the National Academy of Sciences*, 108(5):1821–1826.