Bargaining over Babies: Theory, Evidence, and Policy Implications*

Matthias Doepke
Northwestern University

Fabian Kindermann
Universität Bonn

March 2015

Abstract

It takes a woman and a man to make a baby. This fact suggests that for a birth to take place, the parents should first agree on wanting a child. Using newly available data on fertility preferences and outcomes, we show that, indeed, babies are likely to arrive only if both parents desire one, and there are many couples who disagree on having babies. We then build a bargaining model of fertility choice and match the model to data from a set of European countries with very low fertility rates. The distribution of the burden of child care between mothers and fathers turns out to be a key determinant of fertility. A policy that specifically lowers the child care burden on mothers is more than twice as effective at increasing the fertility rate compared to a general child subsidy.

*We thank seminar participants at the Chicago Fed, Groningen, Northwestern, the SED Annual Meeting, and the 2014 Conference on Demographic Economics at the University of Iowa for comments that helped to substantially improve the paper. Financial support from the National Science Foundation (grant SES-1260961) is gratefully acknowledged. Veronika Selezneva provided excellent research assistance. Doepke: Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208 (e-mail: doepke@northwestern.edu). Kindermann: Institute for Macroeconomics and Econometrics, University of Bonn, Adenauerallee 24–42, 53113 Bonn, Germany (e-mail: fabian.kindermann@uni-bonn.de).

1 Introduction

A basic fact about babies is that it takes both a woman and a man to make one. The fact suggests that some form of agreement between mother and father is required before a birth can take place.[1] In this paper, we present a bargaining model of fertility that builds on this need for agreement, and examine its implications for how different policies, such as child subsidies and public provision of child care, affect childbearing.

The economic literature on fertility choice has mostly relied on unitary models of household decision making. In a unitary model a common objective function for the entire household is assumed to exist, and hence there are no conflicts of interest between spouses. In contrast, in our model husbands and wives have separate preferences and bargain over household decisions. For a birth to take place, both spouses have to prefer an additional child over the status quo.

The key novel implication of this setup is that not just the overall costs and benefits of children matter for fertility (which is what the existing literature focuses on), but also the distribution of costs and benefits within the household. Specifically, if one of the spouses (say, the mother) would end up bearing most of the burden of child care, she may be unwilling to have a child even if her spouse were strongly in favor. In such a setting, a general subsidy for child bearing would be less effective in raising fertility compared to a policy that specifically addresses the intra-household distribution of the burden of child rearing. In addition, if the utility benefit from an additional child differs across spouses, the effectiveness of a policy also rides on whether the costs and benefits of childbearing are properly aligned for each spouse.

An uneven distribution of the burden of child bearing may still not matter for fertility if the spouses can commit to intra-household transfers that compensate the spouse doing most of the work for her efforts. We adopt a household bargaining framework that features limited commitment. Specifically, we employ a version

[1] While exceptions from this rule are possible (such as cases of rape, deception, and accidental pregnancy), they do not account for a major fraction of births in most places and will not be considered here.

of the separate-spheres bargaining model of Lundberg and Pollak (1993b). In a given period, the household allocation is determined through Nash bargaining between husband and wife, where the outside option consists of non-cooperation within the marriage. In the non-cooperation state, the spouses are still married, but do not cooperate with each other, thus forgoing some of the returns to scale accruing to a joint household. The division of labor in this state (such as child care) reverts to a non-cooperative arrangement as well.

While bargaining is efficient within the period, the spouses cannot commit to specific transfers or other actions in future periods. Instead, the allocation in each period is determined through Nash bargaining given the period-specific non-cooperation threat points. This matters for fertility if having a child affects future outside options. In particular, if under the non-cooperative allocation one spouse is stuck with most of the burden of child care, this spouse loses future bargaining power if a birth takes place, and thus may be less likely to agree to having a child.

We confront the predictions of our theory with empirical evidence from the Generations and Gender Programme (GGP), a longitudinal data set covering 19 European countries that includes detailed information on fertility preferences and fertility outcomes. Existing data on fertility preferences has generally taken the form of a preferred total number of children. In contrast, the GGP data contains information on the preferences of both spouses for having another (or a first) child at the present time, i.e., at the time bargaining is taking place. In addition, the survey has a panel structure, so that fertility outcomes for the same couple can be observed in a subsequent period. This data makes it possible to link individual fertility preferences to subsequent outcomes for each parity. The key prediction of our theory is that childbirth should be highly likely only when both potential parents agree on wanting to have a child. The data backs up this prediction. In a regression framework, the importance of agreement can be captured by regressing fertility on each parent's fertility intentions as well as an interaction term, i.e., the product of the parents' intentions. The interaction term captures the additional effect of agreement on the probability of a birth. In all our regressions, the coefficient on the interaction term is significant and much larger than

the coefficients on each spouse's intentions.

To quantify our model, we parameterize the model to match the empirical relationship between fertility intentions and outcomes in the GGP data. The calibration procedure yields an estimate of the underlying preference distribution for child birth, by gender. Interestingly, our estimation reveals a substantial level difference in child preference across genders, with women (on average) valuing childbirth much more than men. In the final step, the quantified model can be used to compare the effectiveness of alternative policies aimed at supporting childbearing. The calibrated model also implies that variation in the share of the burden of childcare that is carried by mothers versus fathers can have a significant impact on the total fertility rate. In particular, the results suggest that one important reason behind the currently extraordinarily low fertility rates in Eastern Europe is that in the low fertility countries, women do a much larger share of child-related work compared to European countries where fertility is higher. We also show that policies intended to raise fertility rates are much more effective when targeted at mothers instead of fathers.

In the following section, we compare our work to the existing literature. In Section 3, we analyze data from the Generations and Gender Programme and document the prevalence of disagreement over fertility among couples, as well as the importance of agreement for a birth to take place. In Section 4, we provide a theoretical analysis of a simplified version of our model, and in Section 5 the full quantitative model is developed. In Section 6 we match the model to the GGP data. Policy simulations are described in Section 7, and Section 8 concludes.

2 Relationship to Existing Literature

Empirical evidence shows that disagreement about fertility choices is commonplace. Westoff (2010) reports that in 17 out of 18 surveyed African countries men desire more children than women do, with an average gap in desired family size of 1.5 and a maximum of 5.6 in Chad.[2] In addition to a level difference in the

[2]Surveys were carried out in 1999 to 2008; see Westoff (2010), Table 6.2.

desired fertility of women and men, there is also evidence of considerable heterogeneity across households, as we show in Section 3 below.

On the theoretical side, our work is related to a large economic literature on fertility choice, which includes the seminal contributions by Becker (1960), Becker and Barro (1988), and Barro and Becker (1989), among others. Most of this literature is cast in terms of the unitary model of the household, where there is a single objective function for the household and potential conflicting preferences among the spouses do not play a role. Here we focus specifically on the contrast with the smaller existing literature that does take bargaining over fertility into account.

A large class of bargaining models of the household are based on cooperative bargaining, where the household achieves an efficient outcome (at least in a static sense). Since the seminal contributions by Chiappori (1988, 1992), the class of models that assume efficiency is also known as the collective model of the household. In the context of the collective model, fertility choice was studied by Blundell, Chiappori, and Meghir (2005). Cherchye, De Rock, and Vermeulen (2012) empirically evaluate a version of the model of Blundell, Chiappori, and Meghir (2005) and find evidence that bargaining power matters for expenditures on children.[3]

In our theoretical framework, decisions are statically efficient, but inefficiencies may arise from dynamic decision making. A similar idea is explored in Iyigun and Walsh (2007), where couples make decisions in two successive periods. In the first period, husband and wife decide non-cooperatively on the effort they put into their own education. In the second period, the couple bargains over the distribution of consumption and leisure as well as the number of children. Iyigun and Walsh (2007) show that women invest more into their education than what would be Pareto efficient, in order to strengthen their bargaining power in marriage.

This reverse impact channel of fertility on the resource allocation is studied by

[3]In Blundell, Chiappori, and Meghir (2005), the different preferences of women and men regarding fertility are taken as primitives. Another approach is to develop microfoundations that provide specific reasons why the incentives for childbearing may differ between women and men. One such paper is Eswaran (2002), who considers a model in which couples decide about the number of children as well as on how much to invest in their children's health.

Rasul (2008). He proposes a model in which husband and wife derive utility from consumption and having children. There is a fixed cost for each child, and the couple has to decide on fertility before the allocation of consumption is determined. In this setting, Rasul (2008) shows that the ability, or lack thereof, of spouses to commit to future actions plays and important role. Under a lack of commitment, after the arrival of a child there is room for renegotiation between the spouses. The threat point in this negotiation is a non-cooperative equilibrium in which the husband still derives utility from the presence of children, but the burden of raising the children falls on the wife. Thus, having children implies a loss of bargaining power for the wife, and given that the wife is assumed to have the final word on the initial fertility decision, the result is inefficiently low fertility. Rasul (2008) confronts the predictions of his model with household data from the Malaysian Family Life Survey, and finds results that suggest a significant lack of commitment in the marital bargaining process.

In terms of emphasizing the importance of bargaining and limited commitment, within the existing literature Rasul (2008) is the most similar to our approach.[4] There are, however, also important differences. First, in the existing models either the wife decides on her own on fertility (Rasul) while taking the impact on future bargaining into account, or fertility choice is part of joint maximization of surplus. In contrast, we envision fertility choice as a step-by-step process (i.e., a separate decision at each parity) where each spouse has veto power, so that agreement between the spouses is essential. This changes the nature of the choice problem, in a way that we show to be consistent with the data. Second, we consider a dynamic model with multiple periods of childbearing, which allows us to distinguish disagreement over the timing of fertility from disagreement over the total number of children. The third important difference is that we use data that allows an empirical evaluation of the link between fertility preferences and

[4] A recent paper along similar lines is Kemnitz and Thum (2014). Dynamic models of fertility choice that also include implications for the marriage market have been developed by Greenwood, Guner, and Knowles (2003), Caucutt, Guner, and Knowles (2002), and Guner and Knowles (2009). Endogenous bargaining power also plays a central role in Echevarria and Merlo (1999) and Basu (2006), although here fertility is exogenous. The extent of commitment within households is analyzed more generally by Mazzocco (2007). Empirical studies of the link between female bargaining power and fertility include Ashraf, Field, and Lee (2014), who suggest that higher female bargaining power leads to lower fertility rates in a developing-country context.

outcomes conditional on the existing number of children and other covariates. This provides a much more detailed view of the bargaining process than was previously available and leads to new insights.

3 Evidence from the Generations and Gender Programme

We use data from the "Generations and Gender Programme" to evaluate the importance of agreement on fertility decisions. The data is collected in 19 mostly European countries, although the questions concerning fertility preferences are available only for a subset of countries. The questions used to determine fertility preferences and agreement or disagreement among spouses are

> Q1: "Do you yourself want to have a/another baby now?"

and

> Q2: "Couples do not always have the same feelings about the number or timing of children. Does your partner/spouse want to have a/another baby now?"

Data for these questions are available for 11 countries in Wave 1 of the survey (which was carried out between 2003 and 2009), with a total of 35,688 responses.

The participants in the study are surveyed again in Wave 2, which takes place three years after the initial interview. So far, Wave 2 data on fertility outcomes is available for five countries, with more to become available in the next two years. The availability of data on fertility outcomes makes it possible to study the link between gender-specific fertility intentions and outcomes in detail. The sample size for each country in each wave is given in Tables 1 and 2.

The data set contains a great deal of other information. Here we focus on data on fertility intentions, labor supply and wages, and the division of child-care tasks

Table 1: Wave 1 sample with questions about fertility preferences

Country	No. of Respondents		
	female	male	Total
Austria	2,149	1,219	3,368
Belgium	1,159	1,058	2,217
Bulgaria	2,691	1,708	4,399
Czech Republic	1,120	1,276	2,396
France	1,640	1,285	2,925
Germany	1,644	1,281	2,925
Lithuania	1,024	1,175	2,199
Norway	2,488	2,446	4,934
Poland	2,211	1,638	3,849
Romania	1,587	1,835	3,422
Russia	1,640	1,414	3,054
Total	19,509	16,179	35,688

Table 2: Wave 2 sample with questions about fertility preferences and observed fertility

Country	No. of Respondents		
	female	male	Total
Bulgaria	1,898	1,190	3,088
Czech Republic	392	254	646
Germany	576	354	930
France	1,099	816	1,915
Total	3,965	2,614	6,579

among the spouses within the household. These are the key variables to evaluate the predictions of our baseline theory.

We now document a set of key facts from the data that inform our economic model, namely:

1. There is a considerable amount of disagreement within couples about whether to have a (or another) baby.
2. Without agreement, few births take place.
3. The extent and nature of disagreement over fertility is closely related to the distribution of childcare duties between women and men.

We now address each of these observations in turn.

There is a considerable amount of disagreement within couples about whether to have a (or another) baby

In order to document how much agreement or disagreement about having a kid prevails in the populations of different countries, we calculate the share of non-agreeing couples within the population of households in which at least one of the partners wants to have a child. Let AGREE denote the fraction of couples where both spouses desire a baby; HE NO denotes the case where the wife/female partner desires a baby, but the husband/male partner does not; and SHE NO means that he desires a baby, but she does not. We now compute the following disagreement shares:

$$\text{DISAGREE MALE} = \frac{\text{HE NO}}{\text{AGREE} + \text{HE NO} + \text{SHE NO}},$$
$$\text{DISAGREE FEMALE} = \frac{\text{SHE NO}}{\text{AGREE} + \text{HE NO} + \text{SHE NO}}.$$

Figure 1 displays the extent of disagreement over fertility across countries, where the total fertility rate for each country is shown in parentheses. In this graph, if all couples in a country were in agreement on fertility (either both want one or both do not), we would get a point at the origin. In a country that is on the 45-degree line men and women would be equally likely to be opposed to having a baby.

Figure 1: Disagreement over having a baby across countries (all couples)

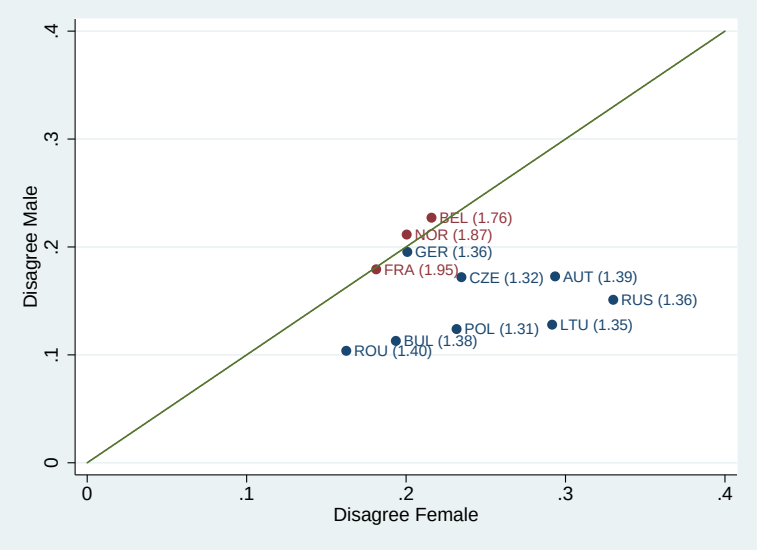


Figure 2: Disagreement over having a baby across countries (couples with no children)

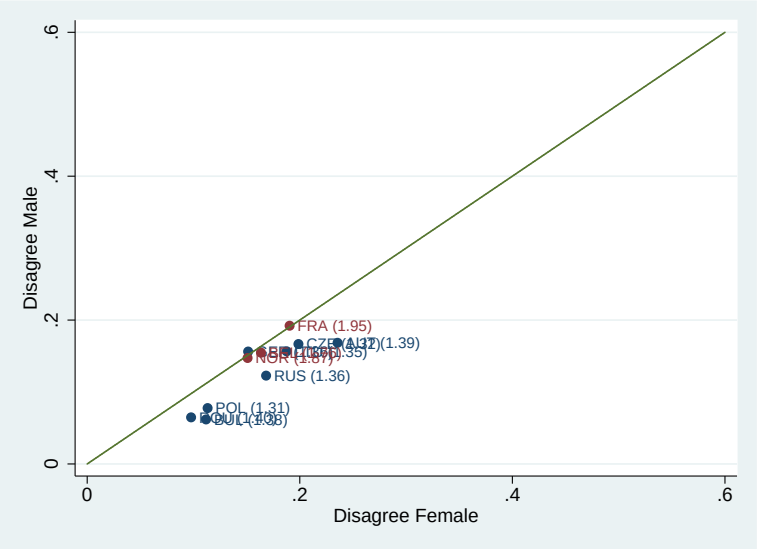


Figure 3: Disagreement over having a baby across countries (couples with one child)

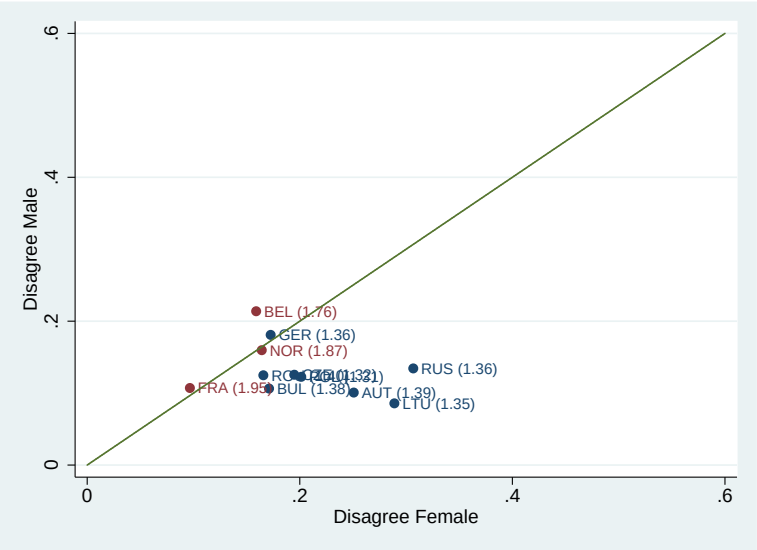
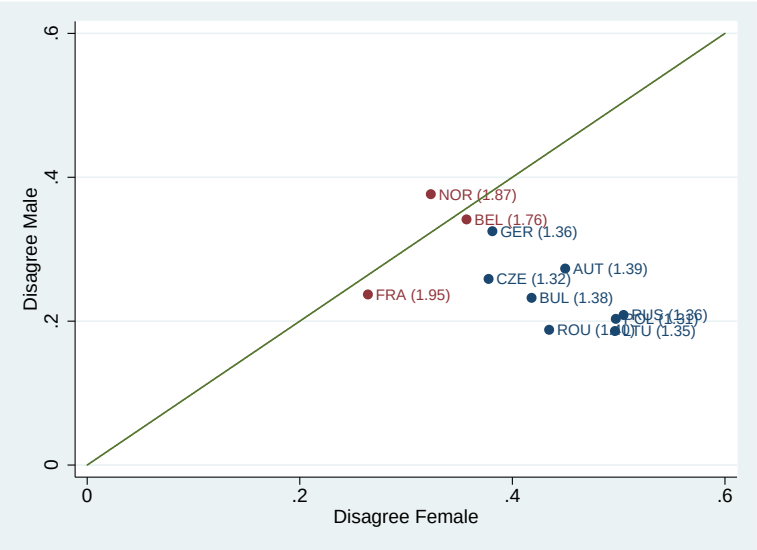


Figure 4: Disagreement over having a baby across countries (couples with two or more children)



The main facts displayed in Figure 1 can be summarized as follows. First, there is a lot of disagreement; in between 25 and 50 percent of couples where at least one partner desires a baby, one of the partners does not (the total disagreement is the sum of the values on the x and y axes). Second, women are more often in disagreement with their partner's desire for a baby than the other way around (i.e., most countries lie to the right of the 45 degree line). Third, the tilt towards more female disagreement is especially pronounced in countries with very low total fertility rates, whereas disagreement is nearly balanced by gender in the countries with a relatively high fertility rate (France, Norway, and Belgium).

The picture as such does not allow conclusions about whether disagreement affects the total number of children a couple ends up with: it is possible that disagreement is about the timing of fertility, rather than on whether or how many children to have overall. This issue will be addressed in the quantitative analysis below by exploiting repeated information on child preferences for couples who took part in both waves of the survey. It is also indicative to consider disagreement exclusively among couples who already have at least two children, which makes it more likely that the potential baby to be born is the marginal child (so that the total number of children would be affected). Figures 2 to 4 break down the data by the number of children already in the family. The main observations here are that among couples who already have at least two children, the extent of disagreement is even larger (50 to 70 percent), and the tilt towards female disagreement in low-fertility countries is even more pronounced.

Without agreement, few births take place

Next, we document that disagreement is an important obstacle to fertility. As already mentioned, analyzing this fully requires taking account of the difference between intentions for the timing of births versus overall fertility intentions, which is described further below. However, the basic facts can be seen through simple regressions of the following form:

$$\text{child}_{t+1} = \beta_0 + \beta_f \cdot \text{fintent}_t + \beta_m \cdot \text{mintent}_t + \beta_{fm} \cdot \text{fintent}_t \times \text{mintent}_t.$$

In this simple linear probability model, fintent$_t \in \{0,1\}$ and mintent$_t \in \{0,1\}$ denote the partners' fertility intentions at time t, and child$_{t+1} \in \{0,1\}$ indicates whether this couple had a child born in the time between the interviews of wave 1 and wave 2. In a world where women decide on fertility on their own, we would expect to find $\beta_f > 0$ and $\beta_m = \beta_{fm} = 0$. If each spouse's intention had an independent influence on the probability of having a baby, we would observe $\beta_f > 0$, $\beta_m > 0$, and $\beta_{fm} = 0$. Finally, if a birth can take place only if the spouses agree on having a baby (i.e., each spouse has veto power), we would get $\beta_f = \beta_m = 0$ and $\beta_{fm} > 0$. In other words, the interaction term is what measures the importance of agreement between the spouses, over and above their general intentions. Least squares estimates for this regression, using pooled data for all available countries, are shown in Table 3.

Table 3: Importance of fertility intentions for outcomes (complete sample)

	Coefficient	Standard Error
fintent	0.1151***	(0.0235)
mintent	0.0614***	(0.0171)
fintent×mintent	0.1736***	(0.0324)

We find that all coefficients are significant, but the interaction term is the largest in size.[5] Perhaps not surprisingly, female intention is more important than male intention. But even for a woman, having her partner agree greatly increases the probability of having a child. A couple that agrees has an almost three times higher likelihood of having a baby than a couple where the husband disagrees, and a more than four times higher likelihood than a couple where the wife disagrees.

The regressions can once again be broken down by parity, i.e., the number of children the couple already has. The same patterns show up at all parities. The need for agreement is most pronounced for couples with no children yet (see Table 4). For these couples, the probability of having a child without agreement is not significantly different from zero.

[5]Formal testing reveals that β_{fm} is at least statistically different from β_m at the 5 percent level.

Table 4: Importance of fertility intentions for outcomes (couples with no kids)

	Coefficient	Standard Error
fintent	0.0418	(0.0147)
mintent	0.0369	(0.0133)
fintent×mintent	0.2096***	(0.0566)

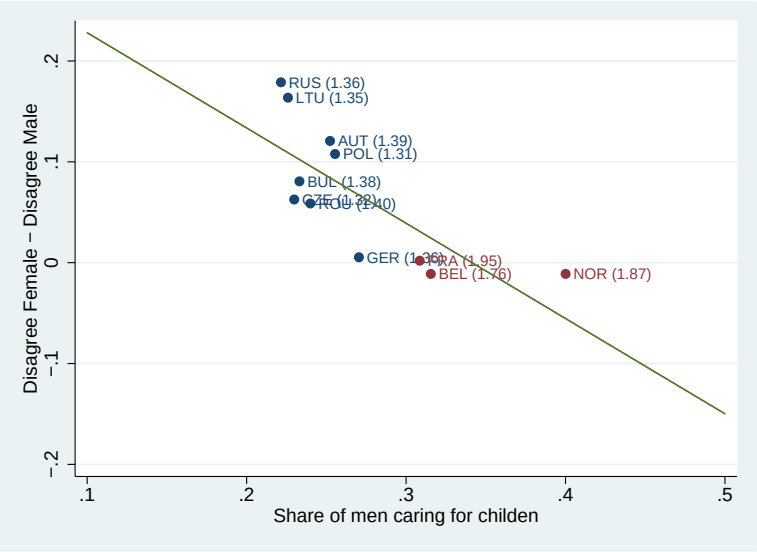
The extent and nature of disagreement over fertility is related to the distribution of childcare duties between women and men

In the theory articulated below, disagreement between spouses regarding fertility can arise because couples cannot commit to a specific allocation of child care duties in the future. This mechanism suggests that at the national level, the extent and nature of disagreement should correlate with policies and attitudes that also affect the distribution of child care duties between mothers and fathers. To show that child care matters in the GGP data, here we calculate the share of men in caring for children by coding the answers to the following questions:

> "I'm going to read out various tasks that have to be done when one lives together with children. Please tell me, who in your household does these tasks?
>
> 1. Dressing the children or seeing that the children are properly dressed;
> 2. Putting the children to bed and/or seeing that they go to bed;
> 3. Staying at home with the children when they are ill;
> 4. Playing with the children and/or taking part in leisure activities with them;
> 5. Helping the children with homework;
> 6. Taking the children to/from school, day care centre, babysitter or leisure activities."

The possible answers to these questions are "always the Respondent", "usually the Respondent", "about equal shares", "usually the Partner" and "always the Partner". We code these answers in 0, 0.1, 0.5, 0.9 and 1 if the respondent is female and 1, 0.9, 0.5, 0.1 and 0 if the respondent is male. We aggregate the answers into one variable by forming a simple mean and calculate the average for every country. This gives us a proxy for the share of men in child care for every country which we can plot against the difference between the share of female disagreement and the share of male disagreement. This yields Figure 5 (which also includes a regression line).

Figure 5: Disagreement over fertility and men's share in caring for children



The figure shows that in countries where women do most of the work in raising children, women are more likely to be opposed to having more children, and fertility is low. While these relationships appear very intuitive and confirm the conventional wisdom on European fertility to some extent, it is important to notice that it takes a particular kind of model to capture these facts. First, a bargaining model is required, as a unitary model is not designed to account for disagreement. Second, the link from disagreement to total fertility suggests that husbands are not able to fully compensate their wives for their child care duties in order to implement their own (higher) fertility preference. We take the

perspective that this is due to limited commitment within the household. The theoretical framework that spells out this mechanism is described next.

4 A Bargaining Model of Fertility

In this section, we develop our bargaining model of fertility choice. We consider the decision problem of a household composed of a wife and a husband. Initially the couple does not have children. To make a child, both partners have to take a joint action, and hence a child is created only if both spouses find it in their interest to participate. Without agreement, the status quo prevails. We start our analysis with the case of a one-time choice of a single child. We contrast the cases of commitment and limited commitment, and show that the distribution of the burden of child care between the spouses is an important determinant of the total fertility rate. Next, we extend the analysis to a two-period model, and show that the evolution of child preferences over time also needs to be taken into account if we want to disentangle the effects of possible policy interventions on period fertility and cohort fertility. These insights lead to the development of a multi-period model with stochastically evolving child preferences in Section 5.

4.1 Commitment versus Limited Commitment in the One-Child Case

Consider an initially childless couple consisting of a wife f and husband w. The couple has to decide on whether to have a child. Each spouse has one unit of time, and the market wages for husband and wife are w_m and w_f. The total cost of a child in terms of consumption is given by ϕ. Utility $u_g(c_g, b)$ of spouse $g \in \{f, m\}$ is given by:

$$u_g(c_g, b) = c_g + bv_g, \tag{1}$$

where $c_g \geq 0$ is consumption, $b \in \{0, 1\}$ indicates whether a child is born, and v_g is the additional utility spouse g receives from having a child compared to the childless status quo.

In addition to the opportunity to have children, an added benefit of marriage is returns to scale in consumption. Specifically, if a married couple cooperates, their effective income increases by a factor of $\alpha > 0$ (or, equivalently, the effective cost of consumption decreases by a factor of $1/(1+\alpha)$). For a cooperating couple, the budget constraint is then given by:

$$c_f + c_m = (1+\alpha)\left(w_f + w_m - \phi b\right). \tag{2}$$

The household reaches decisions through Nash bargaining. Consider first the case of commitment, in which the spouses can commit to a future consumption allocation before having a child. This case amounts to choosing consumption and fertility simultaneously subject to a single outside option.

When there is full commitment, both partners decide simultaneously whether to have a child and how to distribute resources. The decision is made via Nash bargaining, where the outside option is non-cooperation and therefore not having a child, i.e.

$$u_f^d(0) = w_f \quad \text{and} \quad u_m^d(0) = w_m. \tag{3}$$

We denote the ex-post utility of husband and wife (i.e., taking wages, costs of children, and the bargaining outcome into account) as $u_g(0)$ when no child is born and $u_g(1)$ when a child is born, where $g \in \{f, m\}$. We assume equal bargaining weights throughout.[6]

Proposition 1 (Fertility Choice under Commitment)**.** *Under commitment, the couple decides to have a child if the condition:*

$$v_f + v_m \geq \phi(1+\alpha) \tag{4}$$

is met. Moreover, when (4) *is met, we also have:*

$$u_f(1) \geq u_f(0) \quad \textit{and } u_m(1) \geq u_m(0).$$

[6] All results can be generalized to arbitrary weights.

That is, each spouse is individually better off when the child is born. Conversely,

$$v_f + v_m < \phi(1 + \alpha)$$

implies

$$u_f(1) < u_f(0) \quad \textit{and } u_m(1) < u_m(0),$$

i.e., if the couple decides not to have a child, each spouse individually is better off without the child. Taking together, the conditions imply that under commitment the couple always agrees about the fertility choice.

The implication of perfect agreement on fertility among the spouses conflicts with our empirical observation of many couples who disagree on having a child. To allow for disagreement, we now consider a setup with limited commitment.

In this case, bargaining proceeds in two stages. In the first stage, the spouses decide whether to have a child. In the second stage, resources are divided, given the outside option after the fertility decision is sunk. Hence, for each spouse there are two different outside options, for the case where they do have a child and for the case where they don't. This setup captures lack of commitment, in the sense that the spouses are not able to make binding commitments for transfers in the second stage during the first-stage bargaining over fertility (allowing for commitment to such transfers would return us to the full commitment case discussed above).

The outside options conditional on not having children are still given by (3). To formulate the outside options when there is a child, we have to take a stand on who bears the cost of raising children in the non-cooperation state. We assume that cost share of wife and husband are given by χ_f and χ_m with $\chi_f + \chi_m = 1$. The new outside options therefore are:

$$u_f^d(1) = w_f + v_f - \chi_f \phi, \tag{5}$$

$$u_m^d(1) = w_m + v_m - \chi_m \phi. \tag{6}$$

Notice that in the outside option, the spouses still derive utility from the presence of the child. We interpret the outside option as non-cooperation within the

marriage, as in Lundberg and Pollak (1993b). That is, the couple is still married and both spouses still derive utility from the child, but bargaining regarding the allocation of consumption breaks down, and the couple no longer benefits from returns to scale in joint consumption.

We now characterize the fertility choice under lack of commitment.

Proposition 2 (Fertility Choice under Lack of Commitment)**.** *Under lack commitment, we have $u_f(1) \geq u_f(0)$ (the wife would like to have a child) if and only if the condition*

$$v_f \geq \left(\chi_f + \frac{\alpha}{2}\right)\phi \tag{7}$$

is satisfied. We have $u_m(1) \geq u_m(0)$ (the husband would like to have a child) if and only if the condition

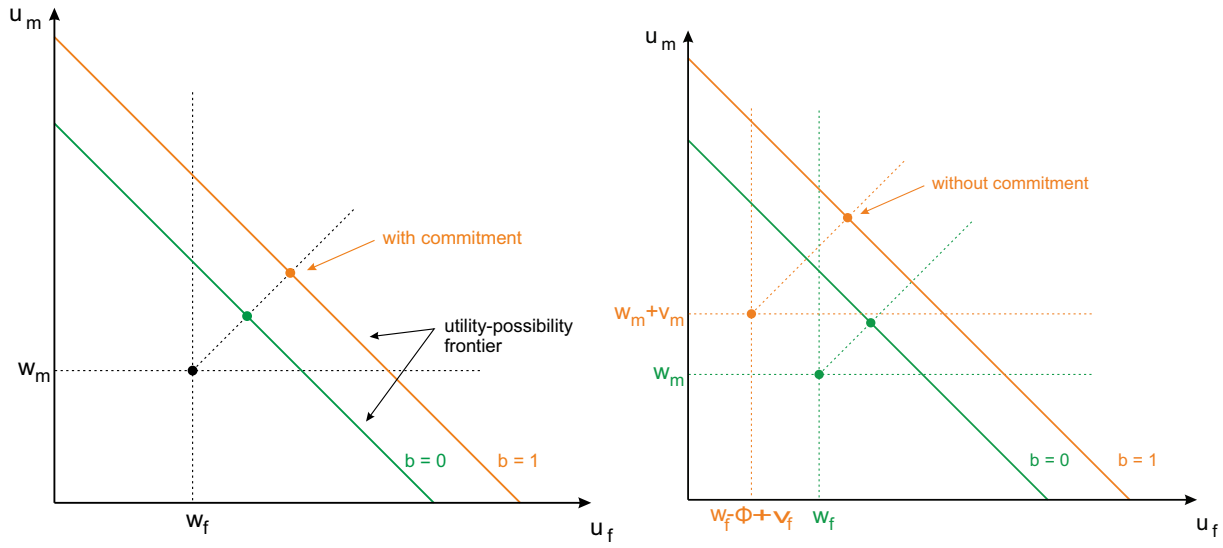
$$v_m \geq \left(\chi_m + \frac{\alpha}{2}\right)\phi \tag{8}$$

is satisfied. The right-hand sides of (7) *and* (8) *are constants. Hence, depending on v_f and v_m, it is possible that neither condition, both conditions, or just one condition is satisfied. Since child birth requires agreement, a child is born only if* (7) *and* (8) *are both met.*

The reason for the possible disagreement is that after the child is born, the outside options of the two partners shift away from the outside options in the full commitment model. Figure 6 illustrates this issue for the case in which the wife bears the entire child cost, i.e., $\chi_f = 1$, $\chi_m = 0$.

The left panel shows the case of full commitment. There the outside option is given by (w_f, w_m). The line $b = 0$ shows the utility possibility frontier for the cases in which the couple does not have children, the line $b = 1$ the frontier for the case of having children. Obviously in this situation the couple wants to have a kid, since it can reach a higher utility level. The utility allocation in this case lies on the intersection between the utility possibility frontier and a 45 degree line starting from the outside option (because of equal bargaining weights). Note that the utility level of having a child will always be higher than the utility level of not having a child, irrespective of the actual distribution of outside options (i.e.

Figure 6: Full versus limited commitment bargaining



the levels of w_f and w_m).[7]

The right panel of Figure 6 shows the case of limited commitment. In this case there are two outside options, the one without children and the one with children. Again, the solution of the bargaining problem is the intersection of the utility possibility frontier and the 45 degree line starting in the outside option. The case drawn here is a case of disagreement. Because she bears all of the child cost and hence loses bargaining power if a child is born, the wife will have a lower utility level in the case with a child compared to the case without children. Hence, she will not agree to a birth and the couple will remain childless, even though they could both be better off with a child if they were able to commit.

[7]Provided that both outside options lie below the utility possibility frontiers, i.e. there actually is a bargaining solution.

4.2 The Distribution of the Burden of Child Care and the Fertility Rate

Our results so far suggest that the distribution of the burden of child care among the spouses matters for fertility; if one spouse bears a disproportionate burden, the spouse will be unlikely to agree to a birth because of the loss in the outside option implied by having a child. We now make this intuition more precise by examining how the average fertility rate in an economy with many couples who are heterogeneous in child preferences depends on the distribution of the burden of child care.

Consider an economy with a continuum of couples. The wages w_f and w_m and the cost shares χ_f and $\chi_m = 1 - \chi_f$ are identical across couples. We interpret the cost parameters as driven partly by government policy, and partly by social norms. For example, there may be a social norm that women do most of the work in raising children, especially in the case of non-cooperation between the couples (which is where the distribution of the burden matters). The extent to which this norm will affect bargaining will also depend on the availability of public child care; if child care can be provided through the market, the husband may be more likely to contribute to the cost of raising children compared to the case where children are always raised within the home by their parents, in which case there would be a greater force towards specialization in child care.

Child preferences are heterogeneous in the population, with a joint cumulative distribution function of $F(v_f, v_m)$. For a child to be born, both (7) and (8) have to be satisfied. For ease of notation, we denote the threshold values for the wife's and husband's child preference above which they would like to have a child by $\tilde{v}_f$ and $\tilde{v}_m$:

$$\tilde{v}_f = (\chi_f + \alpha/2)\,\phi, \tag{9}$$

$$\tilde{v}_m = (\chi_m + \alpha/2)\,\phi = (1 - \chi_f + \alpha/2)\,\phi. \tag{10}$$

The expected number of children $E(b)$ (i.e., the fraction of couples who decide to

have a child) is given by:

$$E(b) = 1 - F\left(\tilde{v}_f, \infty\right) - F\left(\infty, \tilde{v}_m, \right) + F\left(\tilde{v}_f, \tilde{v}_m\right). \tag{11}$$

That is, the couples that don't have a child are those where either the wife's or the husband's fertility preference is below the threshold; the last term is to prevent double counting couples where both spouses are opposed to having a child.

To gain intuition for how fertility depends on the distribution of child costs, it is useful to consider the case of independent distributions $F_f(v_f)$ and $F_m(v_m)$ for female and male child preferences, so that $F(v_f, v_m) = F_f(v_f)F_m(v_m)$. Expected fertility can then be written as:

$$E(b) = 1 - F_f\left(\tilde{v}_f\right) - F_m\left(\tilde{v}_m\right) + F_f\left(\tilde{v}_f\right) F_m\left(\tilde{v}_m\right). \tag{12}$$

If the distribution functions are differentiable at $\tilde{v}_f$ and $\tilde{v}_m$, the marginal effect of a change in the female cost share χ_f on fertility is:

$$\frac{\partial E(b)}{\partial \chi_f} = \phi F_m'\left(\tilde{v}_m\right)\left[1 - F_f\left(\tilde{v}_f\right)\right] - \phi F_f'\left(\tilde{v}_f\right)\left[1 - F_m\left(\tilde{v}_m\right)\right]. \tag{13}$$

The first (positive) term represents the increase in the number of men who agree to have a child if the female cost share χ_f increases (and hence the male cost share declines), and the second (negative) term is the decline in agreement on the part of women. The first term has two components: $F_m'\left(\tilde{v}_m\right)$ is the density of the distribution of male child preferences at the cutoff, which tells us how many men switch from disagreeing to agreeing with having a child as χ_f rises. The second component $1 - F_f\left(\tilde{v}_f\right)$ is the fraction of women who agree to have children. This term appears because the husband switching from disagreeing to agreeing only results in a birth of the wife also agrees. If most women are opposed to having a child, an increase in male agreement has only a small effect on fertility. In the same way, the negative impact if the decline in female agreement on fertility, measured by $F_f'\left(\tilde{v}_f\right)$, is weighted by the share of men agreeing to have a child $[1 - F_m\left(\tilde{v}_m\right)]$.

The terms for the existing fractions of women and men agreeing to have a child in

(13) introduce a force that leads to high fertility if agreement on having children is balanced between the genders. In the extreme, if all women were opposed to having a baby but at least some men wanted one, the only way to raise fertility would be to lower the female cost share (and vice versa if all men were opposed). The overall relationships between cost shares, agreement rates, and fertility can be fully characterized when child preferences are uniform, so that the densities $F'_f(\tilde{v}_f)$ and $F'_m(\tilde{v}_m)$ are constant. In particular, if female and male fertility preferences have the same uniform densities (but potentially different means), fertility is maximized when equal fractions of women and men agree to having a child. If one gender has more concentrated fertility preferences (higher density), fertility is maximized at a point where the rate of agreement in this gender is proportionately higher also. The following proposition summarizes the results.

Proposition 3 (Effect of Distribution of Child Cost on Fertility Rate)**.** *Assume that the female and male child preferences follow independent uniform distributions with means μ_g and densities d_g for $g \in \{f, m\}$. Then expected fertility $E(b)$ is a concave function of the female cost share χ_f, and fertility is maximized at:*

$$\hat{\chi}_f = \max\left\{1, \min\left\{0, \frac{1}{2} + \frac{1}{2\phi}\left[\mu_f - \mu_m + \frac{1}{2}\frac{d_m - d_f}{d_f d_m}\right]\right\}\right\}. \tag{14}$$

Hence, if women and men have the same preferences ($\mu_f = \mu_m$, $d_f = d_m$), fertility is maximized when the burden of child care is equally shared. Moreover, if the distributions of female and male preferences have the same density ($d_f = d_m$), equal shares of men and women agree to having a child at the maximum fertility rate, even if $\mu_f \neq \mu_m$ (provided that $\hat{\chi}_f$ is interior). If $d_f \neq d_m$, at $\hat{\chi}_f$ more individuals of the gender with the more concentrated distribution of preferences (higher d_g) agree to having a child than individuals of the gender with more dispersed preferences. Specifically, fertility is maximized when the ratio of of agreement shares $(1 - F_f(\tilde{v}_f))/(1 - F_m(\tilde{v}_m))$ is equal to the ratio of densities d_f/d_g.

The result suggests that if the distribution of the burden of child care is not at the fertility-maximizing level, the fertility rate could be raised by policies that shift child care responsibilities in a particular direction. Likewise, subsidies for childbearing would be more or less effective depending on whether they specif-

ically target one of the spouses (say, by providing publicly financed alternatives for tasks that previously fell predominantly on one spouse). For a concrete policy analysis, we need to add more structure to the analysis to specify whether child care is in terms of goods spending or time, how taxes are raised, and so on. We do this in Section 5 in a more elaborate quantitative version of our theory. When matched to the GGP data, that model indeed predicts that the effectiveness of policies designed to promote child bearing crucially depend on how the policies are targeted.

For non-uniform distributions of child preferences, the same intuitions regarding the effects of a change in cost shares that arise from Proposition 3 still apply locally. In particular, given (13), the local effect of a change in cost shares is driven by the density of the child preferences of each gender and the existing shares of agreement and disagreement by gender. Global results can only be obtained by placing at least some restrictions on the overall shape of preferences.[8] Empirically, we do not have information on the global shape of child preferences away from the cutoffs, because we only observe a binary variable on child preferences. We therefore use uniform distributions in the quantitative implementation of the dynamic model described below, while noting that the measured effects should be considered to be locally valid. In the quantitative model, we also allow for correlation in child preferences within households. In the mathematical appendix, we show that the results of Proposition 3 also go through in the correlated case.

4.3 The Timing of Births in a Two-Period Model

The analysis so far shows that limited commitment can potentially account for our observations on agreement and disagreement on having children, and that a limited commitment model implies that cultural norms or policy measures that

[8]One can even construct cases (albeit unrealistic ones) where fertility is maximized when one gender bears the entire burden of child care. For example, consider a preference distribution (identical between men and women) where 50 percent of each gender want to have a child even if they have to bear the entire child cost, whereas the other 50 percent agree to having a child only if they bear none of the cost. In this case, 50 percent of couples have a child if one spouse bears all the cost, whereas only 25 percent of couples have a child if both spouses make a contribution.

affect the distribution of the burden of child care within the family can affect fertility outcomes. However, a limitation of the static model is that it does not make a distinction between the timing of births and the total number of births. In a dynamic setting, there is an important distinction between disagreement among spouses about the total number of children they want to have, and disagreement about when to have them. In the extreme, one can envision a setting in which all couples agree on how many children they want to have, and the only source of conflict is whether to have them early or late. In this case, an intervention that reshuffles the burden of child care between the spouses may affect when people have children, but it would not affect the ultimate outcome in terms of the total number of children per couple. If the intention of policy is to raise fertility rates, understanding whether policy affects total fertility or only the timing of fertility is clearly important.

In this section, we extend our analysis to a two-period setting to clarify how the issue of the timing of fertility versus total fertility is related to the persistence of child preferences between periods. In the quantitative model introduced in Section 5 below, this will allow us to use repeated observations of the child preferences of a given couple from multiple waves of the GGP survey to pin down this crucial aspect of the analysis.

As before, there is a continuum of couples, and the wages w_f and w_m, the child cost ϕ, and the cost shares χ_f and $\chi_m = 1 - \chi_f$ are identical across couples and over the two periods $t = 1, 2$. The child cost accrues only in the period when a child is born (to be relaxed in Section 5). Preferences are as in (1), but extending over two periods with discount factor β, where $0 < \beta \leq 1$. Child preferences in the second period may depend on the fertility outcome in the first period. First-period child preferences are denoted as $v_{f,1}$, $v_{m,1}$, and second-period preferences are given by $v_{f,2}$ and $v_{m,2}$. Hence, the expected utility function is:

$$E\left(u_g(c_{g,1}, b_1, c_{g,2}, b_2)\right) = c_{g,1} + b_1 v_{g,1} + \beta E\left(c_{g,2} + b_2 v_{g,2}\right). \tag{15}$$

The expectations operator appears because we allow for the possibility that child preferences in the second period are realized only after decisions are made in the first period. As before, we focus on the case of limited commitment. In each pe-

riod, the spouses bargain ex post over consumption after the fertility decision has been made; in addition, the spouses are unable to commit to a specific consumption allocation in the second period during the first period. There is no savings technology, so that (in the case of cooperation) the per-period budget constraints are as in (2) above. In addition, the outside option of non-cooperation only affects a single period. That is, a non-cooperating couple in the first period returns to cooperation in the second period.

The second period of the two-period model is formally identical to the static model, and Propositions 2 and 3 apply. For a given couple with a given preference draw, let $V_f(0)$ and $V_m(0)$ denote equilibrium second-period expected utilities conditional on no child being born in the first period, and $V_f(1)$ and $V_m(1)$ denote expected utilities if there is a first-period birth. Here the dependence of second-period utility on first-period fertility is solely because preferences in the second period are allowed to depend on the fertility outcome in the first period. We start by characterizing the conditions for births to take place.

Proposition 4 (Conditions for Child Birth in Two-Period Model)**.** *In the second period, a birth takes place ($b_2 = 1$) if and only if the following conditions are satisfied:*

$$v_{f,2} \geq \left(\chi_f + \frac{\alpha}{2}\right)\phi \equiv \tilde{v}_{f,2}, \tag{16}$$

$$v_{m,2} \geq \left(\chi_m + \frac{\alpha}{2}\right)\phi \equiv \tilde{v}_{m,2}. \tag{17}$$

In the first period, a birth takes place ($b_1 = 1$) if and only if the following conditions are met:

$$v_{f,1} \geq \left(\chi_f + \frac{\alpha}{2}\right)\phi + \beta\left(V_f(0) - V_f(1)\right), \tag{18}$$

$$v_{m,1} \geq \left(\chi_m + \frac{\alpha}{2}\right)\phi + \beta\left(V_m(0) - V_m(1)\right). \tag{19}$$

Hence, the main change compared to the static case is that when deciding on fertility in the first period, the spouses also take into account how having a child affects their utility in the second period. Depending on how preferences evolve, this effect could go in either direction. If future preferences are uncertain, there can be an option value of waiting, i.e., a couple may delay having a child in the

hope of a more favorable future preference realization.

We now illustrate how the evolution of child preferences determines whether shifts in the distribution of the burden of child care (say, induced by targeted policies) affect the total number of children (denoted by $n = b_1 + b_2$) or just the timing of fertility. We do this by considering two polar cases. The first one is where first-period fertility does not affect preferences in the second period; instead, fertility preferences are drawn repeatedly from the same distribution. In this scenario, shifts in the cost share affect only total fertility, but not the timing of fertility.

Proposition 5 (Level and Timing of Fertility with Independent Draws)**.** *Assume that in both periods, the female and male child preferences follow independent uniform distributions with identical means μ_g and densities d_g for $g \in \{f, m\}$. Then expected fertility $E(b_1)$ and $E(b_2)$ in the two periods depends on the female cost share χ_f as described in Proposition 3. For any χ_f, we also have $E(b_1) = E(b_2)$, so that total expected lifetime fertility $E(n) = E(b_1) + E(b_2)$ satisfies:*

$$E(n) = 2E(b_1) = 2E(b_2).$$

The timing of fertility, as measured by the ratio $E(b_1)/E(b_2)$, is independent of χ_f.

Next, we consider an opposite polar case where having a child in the first period removes the desire for additional children.

Proposition 6 (Level and Timing of Fertility with Fixed Desire for Children)**.** *Assume that in the first period, the female and male child preferences follow independent uniform distributions with means μ_g and densities d_g for $g \in \{f, m\}$. In the second period, preferences depend on first-period fertility: if $b_1 = 1$, we have $v_{f,2} = v_{m,2} = 0$, and if $b_1 = 0$, we have $v_{f,2} = v_{m,2} = (1 + \alpha/2)\,\phi$. Then the total fertility rate is constant for all $\chi_f \in [0, 1]$:*

$$E(n) = E(b_1) + E(b_2) = 1. \tag{20}$$

Fertility in the first period depends on χ_f as described in Proposition 3 for the transformed

parameters $\tilde{\phi} = (1-\beta)\phi$ *and* $\tilde{\alpha} = (\alpha + 2\beta)/(1-\beta)$. *Given that* $E(n)$ *is constant and:*

$$\frac{E(b_1)}{E(b_2)} = \frac{E(b_1)}{1 - E(b_1)}, \tag{21}$$

the cost share χ_f *affects only the timing, but not the level of fertility.*

Clearly, the proposition captures an extreme case where all individuals eventually want to end up with exactly one child, and the only disagreement is over when that child should be born. But the intuition from this example carries over to the general case where a birth leads to at least some downward shift in future fertility preferences. This is a plausible scenario, because as long as the marginal utility derived from children is diminishing, some such downward shift will be present. If this effect is strong, policies that aim to shift the distribution of the burden of child care may have little impact on the overall fertility rate, even when the data in a given cross section suggests a lot of disagreement over fertility.

To deal with this issue and to allow for a meaningful policy analysis, it is essential to capture how child preferences of a given couple shift over time, and how this depends on child birth. Doing this in a quantitatively plausible manner requires a more elaborate model, which is what we turn to next.

5 A Dynamic Model with Evolving Child Preferences

Our analytical results have shown that to understand the ramifications of disagreement over fertility for potential policy interventions, it is essential to consider a setting that can capture how the fertility preferences of a given couple evolve over time. Hence, we now extend our model to a dynamic setup with stochastically evolving preferences that can be matched to the GGP data.

We are modeling couples that are fertile from period 1 to period $T = 8$. Each model period lasts three years of actual time. The first period corresponds to ages 20–22, the second to 23–25, and so on up to period 8 (ages 41–43). Parents raise their children for $H = 6$ periods (corresponding to 18 years). Hence, after

completing fertility, the couple continues to raise its children until all children have reached adulthood by period $T + H$. Couples start out with zero children and can have up to three children. We denote by b the fertility outcome in a given period, where $b = 1$ in case a child is born in the period and $b = 0$ otherwise. Also, n denotes the total number of children of a couple, where $0 \leq n \leq 3$.

In a given period, a person of gender $g \in \{f, m\}$ derives utility from consumption c_g and fertility $b \in \{0, 1\}$. The utility v_g that a person derives from the arrival of a child is stochastic and evolves over time (to be described below). The individual utility of a household member of gender $g \in \{m, f\}$ at age t is given by the value function:

$$V_g^t(a_1, a_2, a_3, v_f, v_m) = E\left[u(c_g, v_g, b) + \beta V_g^{t+1}(a_1', a_2', a_3', v_f', v_m')\right].$$

Here a_1, a_2 and a_3 denotes the ages of the children at the beginning of the period, v_f and v_m are the child preferences of the two partners, and β is a discount factor that satisfies $0 < \beta < 1$. In writing the value function this way, it is understood that c_g and b are functions of the state variables that are determined through bargaining between the spouses. We have $a_i = 0$ for a potential child that has not been born yet. Since in the model no interesting decisions are made after all children leave the house, we assume that parents die at that point and hence $V_g^{T+H+1} = 0$.

As in Section 4 above, utility is linear in consumption and additively separable in felicity derived from the presence of children. Instantaneous utility is given by:

$$u(c_g, v_g, b) = c_g + v_g \cdot b.$$

Notice that the couple derives utility from a child only in the period when the child is born. However, this is without loss of generality, since only the present value of the added utility of a child matters for the fertility decision.

Children are costly as long as their live with their parents. Given the age distribution of children a_i, we can calculate the total number of children that are living

in the household as:

$$n_h = \sum_i \mathbb{1}(0 < a_i < H) + b,$$

where H is the duration of childhood. The cost of raising n_h children is

$$k(n_h) = \phi \cdot (n_h)^{\psi},$$

with $\phi, \psi > 0$. Depending on the value for ψ, we allow for the possibility of economies or diseconomies of scale. Couples split the cost of children according to the cost shares χ_f and χ_m with $\chi_f + \chi_m = 1$. For now, these cost shares are taken as exogenous.

Couples engage in a cooperative Nash-bargaining game without commitment. Specifically, the spouses cannot commit to future transfers. Bargaining takes place regarding the distribution of consumption within a given period, taking the current number of children and also future utility as given. Both spouses participate in the labor market, with gender-specific wages w_g. Hence, analogous to (5) and (6) in the static model, utility in the outside option is:

$$u(c_g, v_g, b) = w_g - \chi_g k(n_h) + v_g \cdot b,$$

that is, each spouse consumes their own labor income net of the cost of taking care of the children.

In the case of a cooperative bargaining solution, the couple's budget constraint reads

$$c_f + c_m = (1 + \alpha) \left[w_f + w_m - k\left(n_h\right) \right].$$

Here $\alpha > 0$ parameterizes increasing returns to joint consumption that the couple can enjoy if there is cooperation. Assuming equal bargaining weights (which can easily be generalized), the solution to the cooperative bargaining game is the solution to the maximization problem:

$$\max_{c_f, c_m} \Big[c_f - \left(w_f - \chi_f k(n_h)\right) \Big]^{0.5} \Big[c_m - \left(w_m - \chi_m k(n_h)\right) \Big]^{0.5}$$

subject to the above budget constraint. Notice that future utility does not enter here, because the evolution of the state variables is unaffected by the current consumption allocation; hence, the bargaining problem regarding consumption is static. Analogous to (23) and (24) in the proof of Proposition 2, the solution to the maximization problem is:

$$c_f(n_h) = w_f - \chi_f k(n_h) + \frac{\alpha}{2}\left[w_f + w_m - k(n_h)\right],$$
$$c_m(n_h) = w_m - \chi_m k(n_h) + \frac{\alpha}{2}\left[w_f + w_m - k(n_h)\right].$$

That is, each spouse receives its outside option plus a fixed share of the surplus generated by cooperation.

Up to this point, the differences to the setup considered in Section 4 are that we allow for more periods and for a richer structure for the costs of children. These changes lead to a more complicated tradeoff when deciding on fertility, because having a child changes the outside option for as long as the child remains in the household. A spouse with a high cost share will realize that her future bargaining power will decrease if a baby is born, which will make her hesitate to agree. Conversely, a spouse with a low cost share will realize that the loss of bargaining power of the other spouse improves her own future bargaining position, which makes having children attractive over and above the direct utility benefit.

We now introduce two additional modifications that are important for matching the model to the GGP data, namely a more general mapping from fertility intentions into outcomes, and a flexible model of the evolution of child preferences over time.

Regarding fertility, both spouses still form their intentions at the beginning of each period, before bargaining over the consumption allocation takes place. Let $i_g \in \{0, 1\}$ denote the intention of spouse g, where $i_g = 1$ denotes that the spouse

would like to have a baby. Formally, i_g is determined as follows:

$$i_g = I\Big\{ E\left[u(c_g, v_g, 1) + \beta V_g^{t+1}(a_1', a_2', a_3', v_f', v_m')|b = 1\right] \\ - E\left[u(c_g, v_g, 0) + \beta V_g^{t+1}(a_1', a_2', a_3', v_f', v_m')|b = 0\right] \Big\}, \quad (22)$$

where $I(\cdot)$ is the indicator function. (22) expresses that a spouse intends to have a child if having a child increases expected utility. In Section 4, we assumed that having a baby requires agreement, i.e., a child was born ($b = 1$) if and only if $i_f = 1$ and $i_g = 1$. In the GGP data explored in Section 3, we found that agreement between the spouses greatly increases the likelihood of having a baby, but nevertheless some births still occur without perfect agreement. We therefore allow for a general mapping of fertility intentions to outcomes that also depends on the existing number of children. Given fertility intentions and the existing number of children n, the probability of having a baby in a given period is given by a function $\kappa(i_f, i_m, n)$. Later on, we will choose this function to match the observed birth probability for each combination of intent and existing number of children in the GGP data. We take this function as exogenous and don't take a stand on how it is generated; some factors that are likely to play a role are natural fecundity (births are not guaranteed even if the spouses agree), imperfect birth control, and measurement error or change over time in fertility intentions. In Appendix A.3, we show that our results are robust to one particular interpretation of the less-than-perfect match from intentions into outcomes, namely altruism between the spouses.

Regarding child preferences, we saw in Section 4.3 that the persistence of child preferences over time determines the extent to which disagreement over having babies matters for the timing of fertility versus lifetime fertility. To allow for persistence, we model child preferences as follows. The couple starts out with an initial preference draw v_f, v_m from a joint uniform distribution with gender-specific means and correlation ρ between the spouses. If no child is born ($b = 0$), with probability π the couple's fertility preferences are unchanged in the next period. With probability $1 - \pi$, the couple draws new fertility preferences from the same distribution. When a birth takes place ($b = 1$), the couple draws new

fertility preferences, where the mean of the distribution depends on the existing number of children. The dependence of fertility preferences on the number of existing children captures the possibility of declining marginal utility from additional children. This process is formalized as follows. The couple first draws potential new preferences from a joint uniform distribution:

$$\begin{bmatrix} \tilde{v}_f \\ \tilde{v}_m \end{bmatrix} \sim U \left(\begin{bmatrix} \mu_{f,n} \\ \mu_{m,n} \end{bmatrix}, \begin{bmatrix} \sigma_f^2 & \rho\sigma_f\sigma_m \\ \rho\sigma_f\sigma_m & \sigma_m^2 \end{bmatrix}, \right).$$

In the initial period, we have $v_g = \tilde{v}_g$. A new potential draw is then made in each period, and the the law of motion for fertility preferences can then be written as:

$$\begin{bmatrix} v'_f \\ v'_m \end{bmatrix} = \begin{cases} \begin{bmatrix} \tilde{v}_f \\ \tilde{v}_m \end{bmatrix} & \text{with probability } 1 - \pi(1-b) \\ \begin{bmatrix} v_f \\ v_m \end{bmatrix} & \text{with probability } \pi(1-b). \end{cases}$$

Here v'_g denotes fertility preferences in the following period. By matching the evolution of fertility preferences to the GGP data (where fertility preferences for the same couple are observed in repeated waves), we can ensure that the model reproduces the proper mapping from current fertility preferences to long-run fertility outcomes.

6 Matching the Model to Data from the Generations and Gender Programme

We want to calibrate the dynamic model to data from the Generations and Gender Programme. We interpret the data from the various countries as driven by the same structural model, but with potential differences across countries in fertility preferences and the distribution of the burden of child care. One might argue that inherent fertility preferences should be comparable across countries.

However, measured differences in child preferences may reflect in differences in child support policies, the work environment, and other country-specific factors affecting fertility that we do not model explicitly. With this in mind, we use all available data to estimate model parameters that are assumed identical across countries (such as the mapping for fertility intentions into outcomes). In contrast, the burden of child care and fertility preferences are matched to the low-fertility countries in our sample, which display fertility patterns in fertility intentions and fertility rates. Our policy experiments in the following section therefore should be interpreted as being valid for the initial conditions of a low fertility country.

We choose the model parameters in two steps. First, a number of parameters are pinned down individually, by either setting them to standard values or estimating them directly from the data. Second, the remaining parameters, concerning the distribution of child preferences and the evolution of preferences over time, are estimated jointly to match data from the low fertility countries.

6.1 Preset and Individually Estimated Parameters

A number of parameters that are less central to our analysis are set to values that are standard in the literature. First we set the discount factor to $\beta = 0.95$, which corresponds to an interest rate of about five percent. Next, we set the economies of scale in the family to $\alpha = 0.4$, as in Greenwood, Guner, and Knowles (2003). We abstract from economies of scale in childbearing and set $\psi = 1$, that is, all children are equally costly.[9] The remaining question is about the (annual) costs a couple incurs for raising a child. The final preset parameter is the level of the child cost ϕ. Given that utility is linear in consumption, ϕ is a scale parameter that does not matter for any of our results regarding fertility.[10] However, for interpreting policy experiments such as child subsidies it is still useful to attach a specific value to ϕ. In reality, child costs are a combination of direct expenses, payments for

[9]This assumption is less restrictive than it seems, because the means of the child preferences for second and the third child would adjust in our estimation according to the assumed economies of scale.

[10]Specifically, if ϕ is increased by a given amount, the the distribution of child preferences can be shifted up by a corresponding amount to give identical results for fertility intentions and outcomes.

child care, forgone earnings, and opportunity costs of reduced leisure. While this makes it difficult to pin down a precise number, the literature suggests a plausible range of the costs. Guner, Kaygusuz, and Ventura (2014) estimate the average annual expenditure on child care to range between $4,851 and $6,414 per year, depending on the age of a child. Adda, Dustmann, and Stevens (2015) quantify the cost of having a child and working for a women to range between €12.6 and €31.1 per day. With about 250 working days per year this leads to a cost between €3,150 and €7,775. Baudin, de la Croix, and Gobbi (2015) estimate the time cost of having a child at 20 to 30 percent of the time endowment of a woman. With an average salary of around €36,000 for full-time working women in Germany, this would imply a cost of r€7,200 to €10,800. In addition, the OECD consumption equivalence scale quantifies the consumption cost of a child to be around 0.3 times the consumption of an adult. Adda, Dustmann, and Stevens (2015) estimate this equivalence scale to be 0.4. The statistical office of Germany estimates the consumption expenditure of couples with children to average at €38,000 in 2011. Using the OECD equivalence scale for a couple with two children, this would lead to an annual expenditure of around €5,000 per year. To reflect all these cost components in our model, we assume that the annual cost of a child amounts to €10,000.

The first parameters that we estimate directly from the data are the probabilities of having a child within three years conditional on the intentions of the male and the female spouse $\kappa(i_f, i_m, n)$. We assume that these parameters do not vary over countries, so we construct them from our full sample for which we have two waves of data, allowing us to link intentions and outcomes (Bulgaria, Czech Republic, France, and Germany).[11] For these countries, we estimate the regression

$$\text{child}_{t+1} = \beta_0 + \beta_m \cdot \text{SHENO}_t + \beta_f \cdot \text{HENO}_t + \beta_{fm} \cdot \text{AGREE}_t.$$

separately for couples with zero, one, and two existing children. β_0 then is the (3 year) baseline fertility rate for all couples, $\beta_0 + \beta_f$ the fertility rate of couples in which only the woman wants a baby, $\beta_0 + \beta_m$ is the fertility rate if only the

[11] We use all available data because the number of data in each cell would become too small if we estimated the regressions separately by country.

man wants a baby, and $\beta_0 + \beta_{fm}$ is the fertility rate for spouses who both want a baby. The results are shown in Table 5. From these regression results, we derive

Table 5: Regressions of child birth on fertility intentions

	$n = 0$	$n = 1$	$n = 2$
Constant (β_0)	0.1243***	0.1087***	0.0334***
	(0.0194)	(0.0259)	(0.0033)
She no (β_m)	0.0301	0.0195	0.0241
	(0.0369)	(0.0111)	(0.0223)
He no (β_f)	0.0261	0.1602***	0.0819**
	(0.0418)	(0.0319)	(0.0323)
Agree (β_{fm})	0.2658***	0.3252***	0.3397***
	(0.0291)	(0.0524)	(0.0380)

the numbers shown in Table 6 for the probabilities of child birth $\kappa(i_f, i_m, n)$. Note that we use a value of zero where the coefficients are not significantly different from zero, although using the point estimate instead does substantially alter our findings.

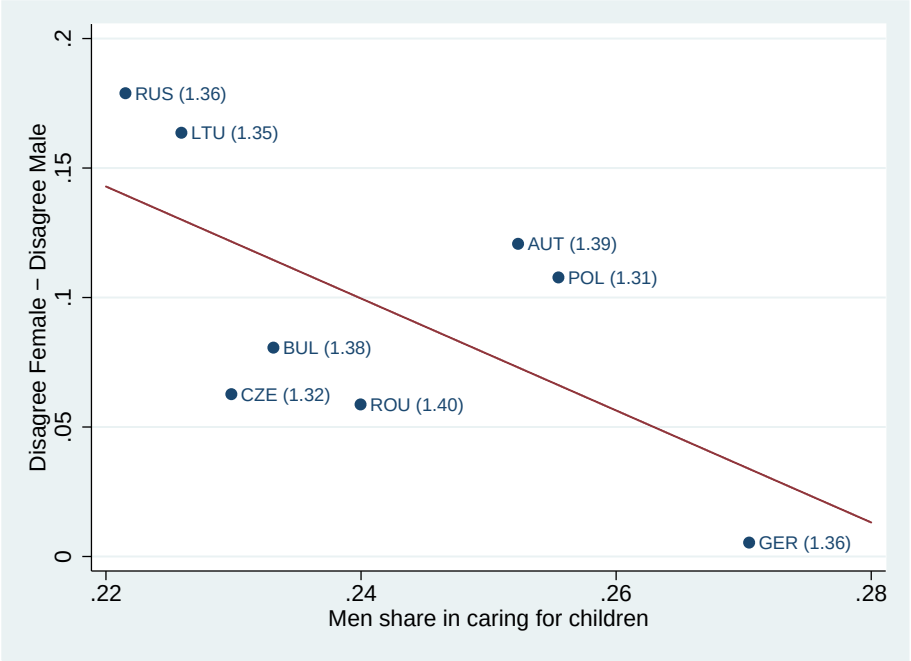
Table 6: Fertility rates by fertility intention (percent of couples with each combination of female intent, male intent, and existing number of children that will have a baby within three years)

Existing children	$n = 0$		$n = 1$		$n = 0$	
	He no	He yes	He no	He yes	He no	He yes
She no	12.43	12.43	10.87	10.87	3.34	3.34
She yes	12.43	39.01	26.89	43.39	11.53	37.31

Next, we pin down the burden of raising a child χ_g. As already shown in Section 3, the Generations and Gender Programme asks individuals which parent carries out specific child care tasks. From these questions, we construct the share of men in total child care (see Figure 5). Figure 7 plots the relevant data from the

low fertility countries that we use in the calibration procedure. We set the male

Figure 7: Male participation in child costs



cost share χ_m to the mean of the share of men in child care for the low fertility countries, which is 0.24. Below, we will also use information on the variation in male cost shares across low fertility countries (which vary between 0.22 and 0.27 in our sample) as target moments to pin down additional parameters.

6.2 Jointly Estimated Parameters

The remaining parameters to be determined concern the distribution of female and male child preference and the persistence of child preferences over time. These parameters are summarized in the following vector:

$$\theta = \begin{bmatrix} \mu_{f,1} & \mu_{f,2} & \mu_{f,3} & \sigma_f & \mu_{m,1} & \mu_{m,2} & \mu_{m,3} & \sigma_m & \rho & \pi \end{bmatrix}'.$$

They include the means of preferences for the first, second, and third child for women and men as well as their standard deviations and within-couple correlation. In addition, we have to determine the persistence of child preferences over time π. To pin down all these parameters we use the following identification strategy.

Means and correlation To pin down the means and the correlation of the distribution of child preferences, we use the reported data on fertility intentions by the two spouses conditional on the number of existing children. Given that fertility can be at most two int the model, for fertility intentions given $n = 2$ we group all couples with two or more children. We generate this data from a pooled sample of the low fertility countries in the Generations and Gender Programme, which are Austria, Bulgaria, the Czech Republic, Germany, Lithuania, Poland, Romania and Russia. We have a total of 25,612 observations with 5,084, 7,664, and 12,864 observations in the $n = 0$, $n = 1$, and $n = 2$ groups, respectively. To pool the sample, we calculate the country specific cross tables of fertility intentions of men and women, using the sample weights. We then take the non-weighted average across countries to derive the pooled intention tables. The results are shown in the first part of Table 7. These 12 data points determine seven target moments, namely six mean parameters and one correlation parameter.

Table 7: Distribution of fertility intentions

		$n = 0$		$n = 1$		$n = 0$	
		He no	He yes	He no	He yes	He no	He yes
Data	She no	50.74	7.40	63.10	8.19	89.27	4.81
	She yes	5.64	36.22	4.55	24.16	2.57	3.35
Model	She no	50.27	7.44	65.15	7.93	87.13	5.99
	She yes	5.75	36.53	3.24	23.67	2.97	3.90

Persistence In order to calibrate the preference persistence parameter π, we use data from all low fertility countries for which we have two waves, namely Bulgaria, the Czech Republic, and Germany. In these countries we select couples that don't have a baby in either wave. We also drop couples in which the female spouse is beyond the age of 35 in the first wave and couples who report that it is physically impossible for them to have a baby. We then look at the 889 couples for which both partners don't want to have a baby in the first wave and calculate their fertility intentions in the second wave. The result is shown in Table 8. This

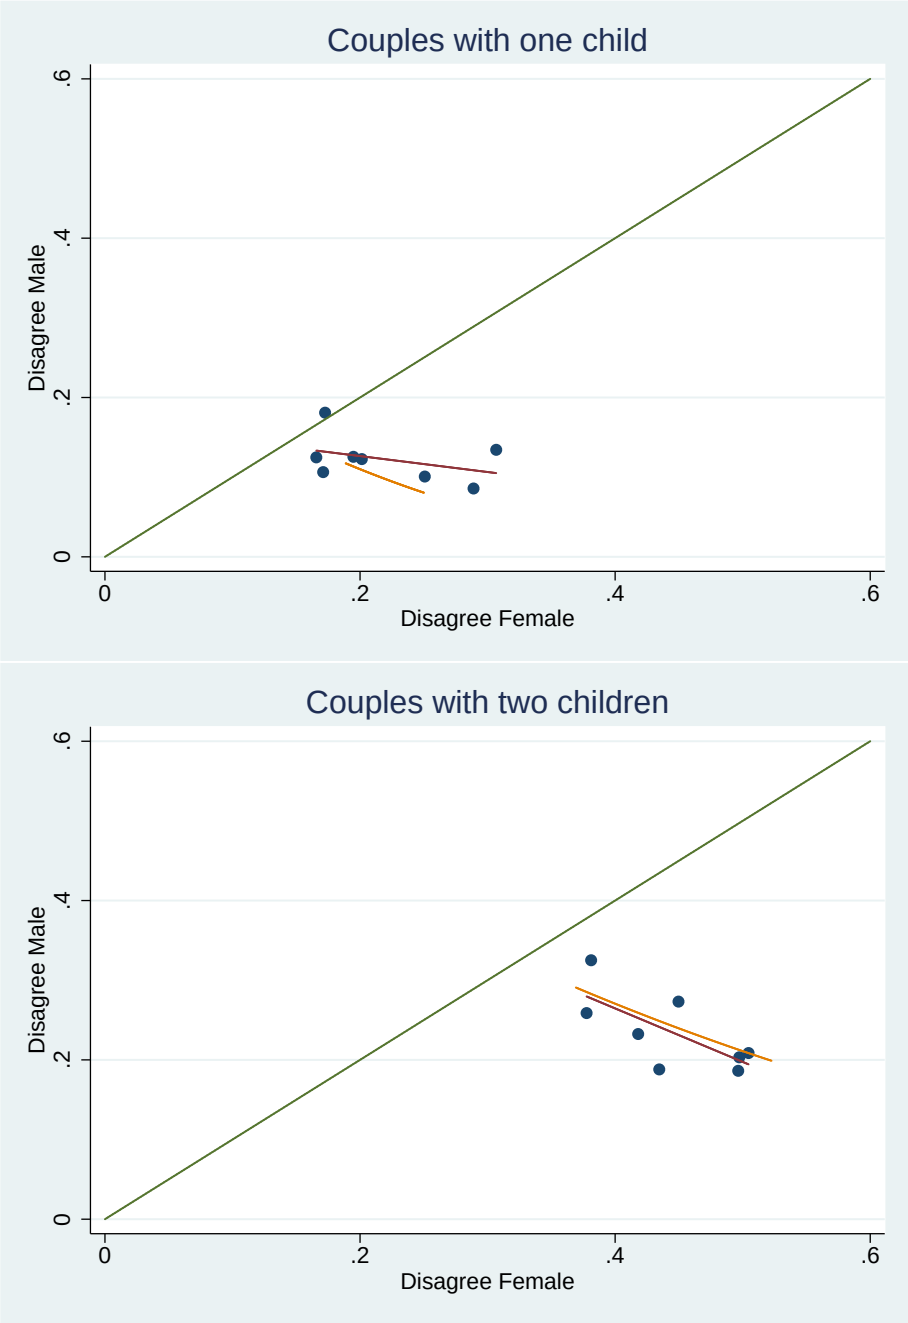
Table 8: Distribution of fertility intentions in the second wave

	Data		*Model*	
	He no	He yes	He no	He yes
She no	85.64	4.42	87.68	3.50
She yes	2.58	7.36	3.00	5.82

table indicates how persistent child preferences are over time. We use the first entry of this table to identify our persistence parameter π.

Standard deviations The last two parameters that we have to pin down are the standard deviations of child preferences σ_f and σ_m. These standard deviations determine how strongly men and women react to changes in the cost of children. Intuitively, if the standard deviation is small, the density of preferences around the cutoff between wanting and wanting a child is high. A small change in child costs will then change the fertility intentions of many individuals, leading to a large change in the fertility rate. The standard deviations are therefore important determinants of the effectiveness of policies aimed at raising fertility. We cannot identify the standard deviations from the distribution of child preferences in Table 7 alone (for the same reason that standard deviations are fixed in a probit model). Instead, we make use of the cross country variation in our sample, namely the relationship between the cost share of men χ_m and relative male and

Figure 8: Fertility intentions across countries



female fertility intentions displayed in Figure 7. Specifically, we look at the distributions of fertility intentions in low fertility countries for couples with one child and couples with two or more children. These distributions are shown in Figure 8. The red lines in the graphs indicate regression lines of the form

$$\text{Disagree male}_i = \beta_0 + \beta_1 \cdot \text{Disagree Female}_i,$$

with i denoting the country index. The outcomes of these regressions are shown in Table 9. To pin down the standard deviations σ_f and σ_m we assume that the

Table 9: Cross country regressions for fertility intentions

	$n = 1$	$n = 2$
β_0	0.1664**	0.5331***
	(0.0489)	(0.1091)
β_1	-0.2001	-0.6715***
	(0.2141)	(0.2275)

differences we observe in fertility intentions are caused by varying cost shares of males in the burden of children across the countries. We therefore let χ_m vary from its benchmark value of 0.24 to the extremes we reported above, i.e. 0.27 and 0.22. The model variants that we simulate with these extremes should then predict the left and the right endpoint of the regression lines in Figure 8, respectively. The model-generated relationships generated by the calibrated model are shown in the figure in yellow.

6.3 Parameter Choices and Model Fit

Let Y denote the 17 target moments we describe above, i.e. the 12 values for the distribution of fertility intentions, the persistence of child preferences, as well as the four end points of the regression lines in Figure 8. In addition, let $\hat{Y}(\theta)$ denote the model simulated counterparts for a set of parameters θ. To pin down

the parameters, we numerically solve the problem

$$\min_{\theta} \quad \left[\hat{Y}(\theta) - Y\right]' \cdot \left[\hat{Y}(\theta) - Y\right],$$

i.e. we minimize a simple residual sum of squares. The resulting set of parameters is shown in Table 10. The model predicted distributions of fertility intentions

Table 10: Calibrated parameters

Description	*Parameter*	*Value*
Mean women first child	$\mu_{f,1}$	186,733
Mean women second child	$\mu_{f,2}$	125,366
Mean women third child	$\mu_{f,3}$	27,624
Std. dev. women	σ_f	86,301
Mean men first child	$\mu_{m,1}$	198,283
Mean men second child	$\mu_{m,2}$	1,680
Mean men third child	$\mu_{m,3}$	-401,956
Std. dev. men	σ_m	347,154
Correlation	ρ	0.9396
Persistence	π	0.8541

as well as the predictions about the persistence of child preferences are shown in Tables 7 and 8. The cross country predictions of fertility intentions are shown as orange lines in Figure 8.

The calibrated model provides a good fit for the data on fertility intentions and the persistence of child preferences over time. In terms of the cross-country variation in fertility preferences, the model makes a good prediction for couples that already have two children. For couples with one child only, the model underpredicts the variation in the intentions of women. Given that for matching overall fertility providing a good fit at the marginal (i.e., last) child is crucial, this should not be of much concern.

The estimated parameters suggest steeply declining marginal utility from hav-

ing children, especially for men. From the second child onwards, women are estimated to have stronger child preferences than men. Intuitively, this occurs because the estimated cost share implies that women carry most of the burden of child care, yet there are still at least some women who desire a second or third child. The estimated rationalizes this by assigning a stronger child preference to women. Child preferences also turn out to be highly persistent and strongly correlated within couples. As argued above, a high persistence of preferences is important for shaping how disagreement versus agreement on children translates into lifetime fertility rates. The high correlation may appear surprising, given that we document substantial disagreement among couples regarding having children. However, at all parities the majority of couples agree that they don't want to have a child, which the model accounts for by highly correlated preferences. The less-than-perfect correlation still leaves enough room for disagreement to arise for a substantial portion of couples.

Table 11: Demographic variables

Total fertility rate	1.4356
Fraction of couples without children	0.1767
Fraction of couples with one child	0.3859
Fraction of couples with two children	0.3899
Fraction of couples with more than two children	0.0474

Table 11 reports some demographic statistics for the model. The model predicts a total fertility rate of the low fertility countries of 1.44, which is a little larger than the average in these countries of 1.36. Given that this number was not targeted, the close fit suggests that the measured fertility intentions translate into overall outcomes in an accurate manner. The model also predicts that after having completed the fertile period, i.e. at the age of 45, the majority of couples has one or two children. Only a small fraction has three children, and about 18 percent of couples are childless. For comparison, the German Statistical Office reports that in 2008, about 19 percent of women between the ages 40 and 49 had no children.

7 Policy Experiments: The Effectiveness of Targeted Child Subsidies

We now turn to the policy implications of our analysis. In many countries, historically low fertility rates are considered a major challenge for future economic prospects, because it is difficult to sustain economic growth with a shrinking population and maintain social insurance systems in an aging population. Already, child bearing is subsidized and publicly supported in various ways in many countries, but there are doubts about the effectiveness of such policies. Here, we study the effect of targeted subsidies in the context of our calibrated model. We assume, in line with Lundberg and Pollak (1993a) and Lundberg, Pollak, and Wales (1997), that subsidies for children can be targeted towards a specific spouse. Intuitively, consider a country where, for mothers, the main component of the burden of child care is forgone earnings, because of an absence of market-based child care and hence the necessity to stay home with the children. In such a setting, public provision of child care centers that allow mothers to go back to work could be considered a policy that is targeted at mothers, whereas a monetary transfer send to the husband would be a policy that is targeted at fathers. Hence, while in the context of the model we speak of monetary transfers, these policies can be interpreted more generally as interventions that specifically relieve the child care burden of one of the spouses.

Formally, let $s_g(n_h)$ denote the total amount of subsidy paid to the partner g for the n_h children currently living in the household. Then the distribution of consumption taking into account subsidies reads

$$c_f(n_h) = w_f - \chi_f k(n_h) + s_f(n_h) + \frac{\alpha}{2}\left[w_f + w_m - k(n_h) + s_f(n_h) + s_m(n_h)\right],$$
$$c_m(n_h) = w_m - \chi_m k(n_h) + s_m(n_h) + \frac{\alpha}{2}\left[w_f + w_m - k(n_h) + s_f(n_h) + s_m(n_h)\right].$$

We now carry out the following experiment. We assume that the government wants to increase the total fertility rate by 0.1 (i.e., one in ten women should have an additional child, increasing the fertility rate from 1.44 to 1.54). It can use subsidies to either women or men to do so. In addition, it can choose to only

pay subsidies for higher-order children, i.e. from the second or the third child onwards.

Figure 9 shows the subsidy amounts that would be needed to increase the total fertility rate by up to 0.1. There are two things to note here. First, whether subsidies are paid for all children or from the second child onwards does not change the amount very much. However, when given for the third child only, the government needs to pay substantially more per child. While for women the annual subsidy needed to increase the total fertility rate by 0.1 is around €2,000 in the former case, it amounts to €7,000 in the latter.

The second and most important feature is that it is much more effective to target subsidies towards women than towards men. Specifically, the subsidy needs to be about 2.5 times larger when targeted towards men than towards women. This finding is novel to our analysis and would not arise in a model that abstracts from bargaining. The reason for the finding is threefold. First, as displayed in Figures 1 to 4, in the low fertility countries that we calibrate to many more women are opposed to having another child than men. Thus, women are more likely to be pivotal in the household decision (see Proposition 3), which means that subsidies directed to women are more effective. There are also additional forces that amplified this effect. The second reason for subsidies to women being more effective is related to the distribution of fertility preferences. Looking at the estimation results in Table 10, we can see that the variance of child preferences for women is lower than for men, indicating that there are more women close to the threshold at which they switch to wanting a baby. Consequently, with a given subsidy the government can incentivize more women than men to switch their opinion towards having another baby. The third reason can be gleaned from the fertility rate regressions in Table 5, where we can see that women have a larger impact on the fertility decision in the household. In fact, the coefficient regarding the male fertility intention on the fertility outcome of the family is never significantly different from zero. These three reasons combined together imply that subsidies that are targeted towards women are much more likely to succeed in raising the total fertility rate.

The data shown in Figure 9 does not allow us to compare the desirability of sub-

Figure 9: Annual subsidy amounts to men and women

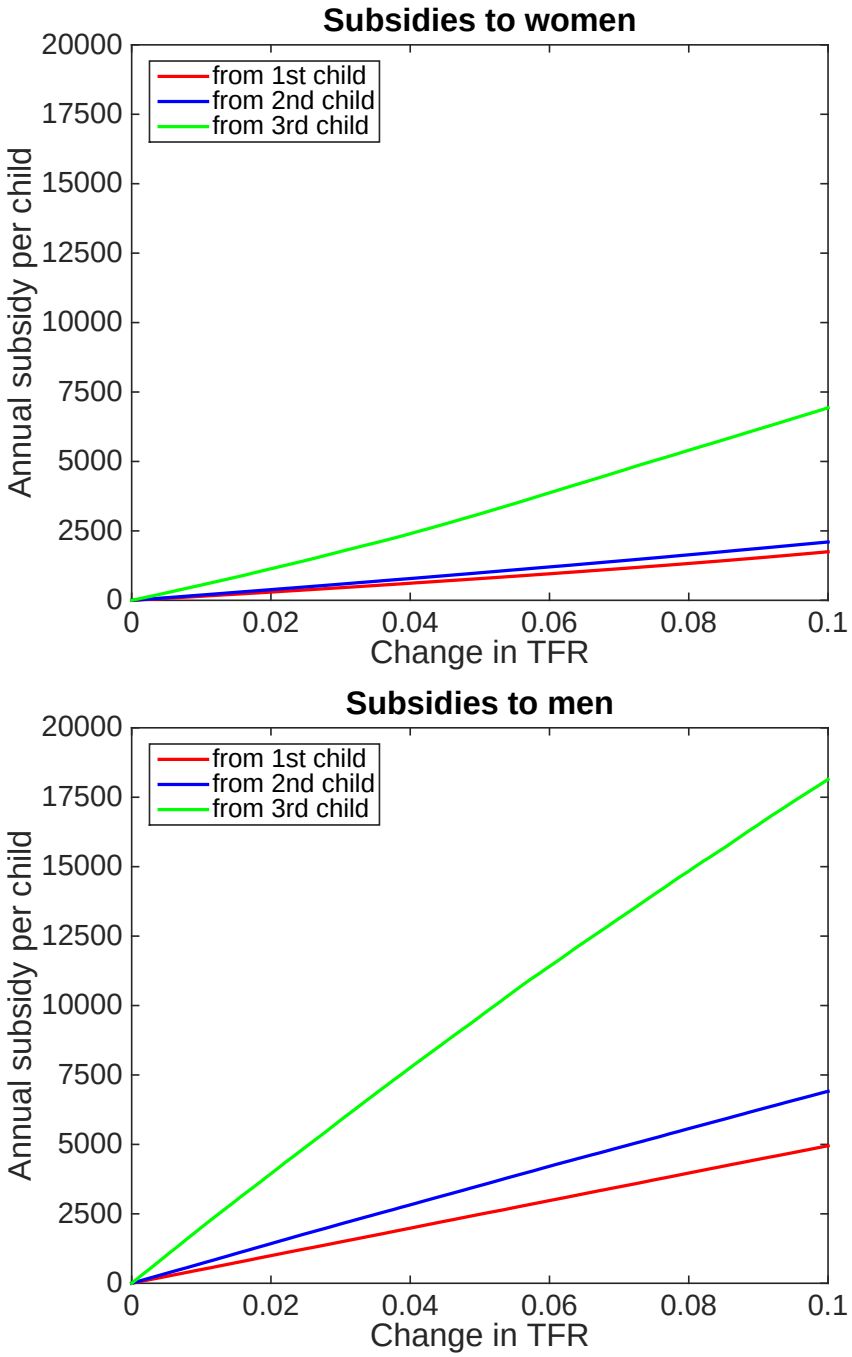
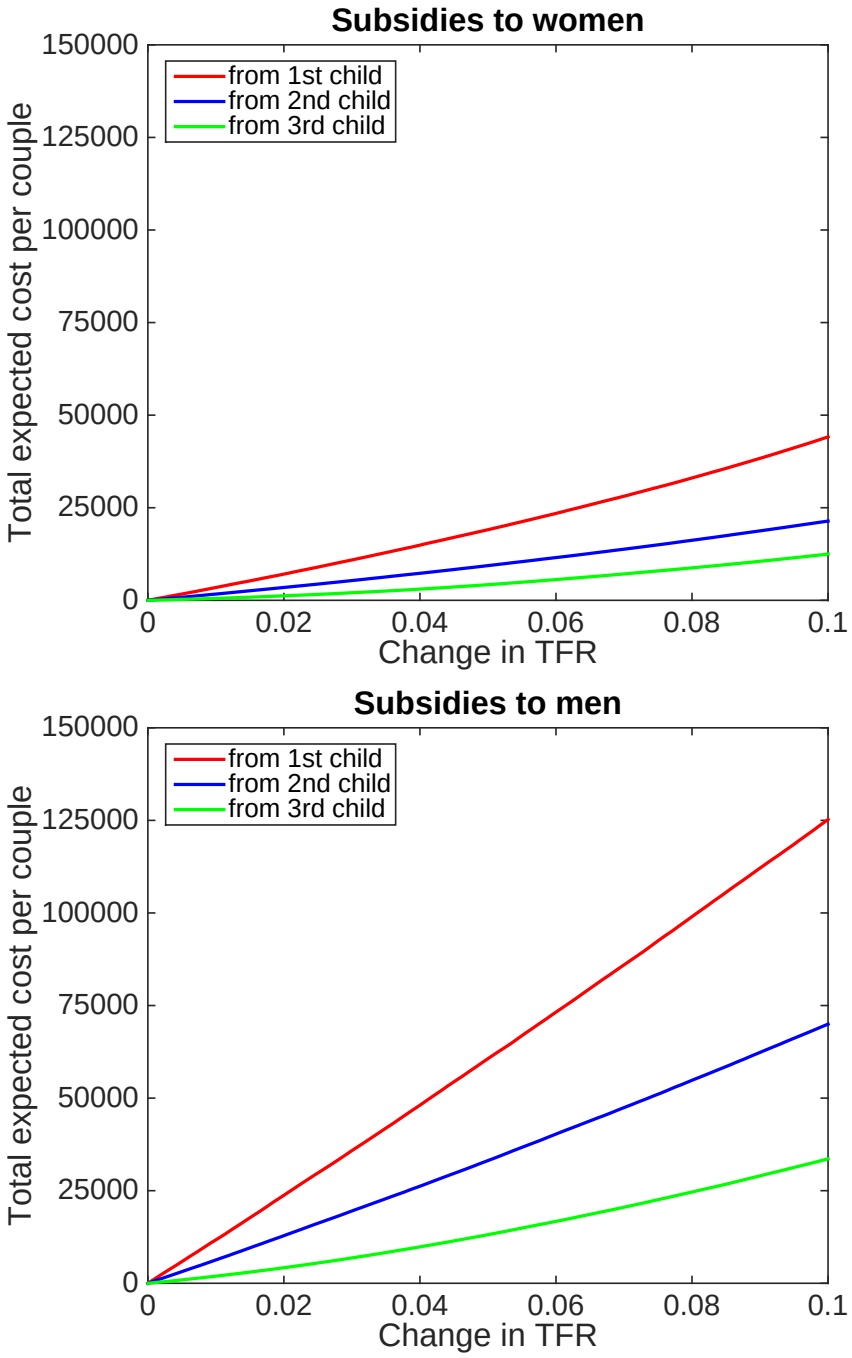


Figure 10: Average cost per couple



sidies that target all children versus, say, only third children and up. While the per-child subsidy needs to be higher when only higher-order births are subsidized, there are also fewer of those children. The total cost of each version of the subsidy is summarized in Figure 10, which displays the average cost per couple over their whole life course that needs to be paid by the government to raise fertility a given amount. The figure reveals that while the required per-child subsidy is the smallest if given for all children, the total cost of this policy is in fact the largest. Increasing the total fertility rate by 0.1 can be achieved with an average lifetime subsidy of €20,000 per couple, but the policy is about twice as expensive if all children are subsidized. This finding can be explained by the distribution of completed fertility in Table 11. The table shows that there are many couples who would have at least one child even without the subsidy. All subsidies given to these couples for the first child do not affect the total fertility rate. These sunk costs make the policy costly in the aggregate. Targeting subsidies to higher order children is more cost effective, since the program is better targeted towards marginal children.

Summing up, we can draw two conclusions from this analysis. First, subsidies are most effective if targeted at higher-order children, which are more likely to be marginal. Second, subsidies should be targeted towards women, who are much more likely to be pivotal for a couple's fertility decision. Raising fertility by subsidizing men is about 2.5 times more costly compared to subsidizing women. Hence, at least in the low fertility environment that our model is calibrated to, accounting for the patterns of agreement and disagreement on having babies makes a huge difference for policy effectiveness.

8 Conclusions

In this paper, we have examined the demographic and economic implications of the simple fact that it takes agreement between a woman and man to make a baby. Using newly available data from the Generations and Gender Programme, we have shown that disagreement between spouses about having babies is not just a theoretical possibility, but commonplace: for higher-parity births, there are

more couples who disagree about having a baby than couples who both want one. We have also shown that disagreement matters for outcomes, in the sense that a baby is unlikely to be born unless both parents desire one. We interpret the data using a model of marital bargaining under limited commitment, and show that our calibrated model provides a close match for the data on fertility intentions and outcomes.

Our findings have both positive and normative implications for the economics of fertility choice. On the positive side, our theory suggests a novel determinant of a country's average fertility rate, namely the distribution of the burden of child-care between mothers and fathers. If one gender carries most of the burden, we would expect to observe a lopsided distribution of fertility intentions, and the fertility rate can be low even if childbearing is highly subsidized overall. Indeed, in the sample of European countries in the GGP data, we find that all low fertility countries are characterized by many more women than men being opposed to having another child.

In terms of normative implications, the analysis suggests that policies that aim at raising the fertility rate will be more effective if they specifically target the gender more likely to disagree with having another child. In our quantitative model calibrated to the European low fertility countries, we find that a child subsidy that specifically lowers women's burden in child care (i.e., by publicly funding child care that allows a mother to return to work earlier) is, dollar for dollar, more than twice as effective at raising fertility than a subsidy targeted at fathers. In many industrialized countries, the currently extremely low fertility rates are projected to lead to major problems in sustaining social insurance system in the future, so that raising fertility is a key policy challenge. We believe that examining policies from the perspective of their effect on agreement and disagreement within couples on fertility is an important direction for future theoretical and applied research. One immediate implication is that optimal policy will be country specific, because patterns of disagreement over fertility vary widely across countries. In the GGP sample, it is notable that the high fertility countries (Belgium, France, and Norway) already have broadly balanced fertility intentions between women and men, so that there is less need for targeted policies.

We have left some aspects of our analysis deliberately simplified to focus on the core issue of fertility intentions and outcomes in a setting with bargaining under limited commitment. To further refine the policy implications, in the next step of this research program it will be necessary to add detail to other aspects of the theory. In particular, here we do not take a stand on exactly what the burden of child care consists of. For policy implications, it is important to know whether, say, a parent's ability to return to work, monetary expenses, or the division of general household chores are the crucial issue leading to disagreement. We plan to examine such dimensions both empirically and theoretically in future research.

A Mathematical Appendix

A.1 Proofs for Propositions

Proof of Proposition 1: The bargaining problem can be solved via backward induction, i.e., we first solve for the ex-post allocation for a given fertility choice, and then consider the optimal fertility choice in the first stage.

If the couple decides not to have a child ($b = 0$), then resource allocation is determined by the maximization problem:

$$\max_{c_f, c_m} \quad [c_f - w_f]^{0.5} [c_m - w_m]^{0.5} \quad \text{s.t.} \quad c_f + c_m = (1 + \alpha) [w_f + w_m] .$$

Here α is an efficiency scale factor that defines the surplus of a marriage. Individual consumption in this case is given by:

$$c_f(0) = w_f + \frac{\alpha}{2} [w_f + w_m] \quad \text{and} \quad c_m(0) = w_m + \frac{\alpha}{2} [w_f + w_m] ,$$

and utilities are:

$$u_f(0) = w_f + \frac{\alpha}{2} [w_f + w_m] \quad \text{and} \quad u_m(0) = w_m + \frac{\alpha}{2} [w_f + w_m] .$$

If the partners do decide to have a child ($b = 1$), the resource allocation solves the maxi-

mization problem:

$$\max_{c_f,c_m} \quad [c_f + v_f - w_f]^{0.5} [c_m + v_m - w_m]^{0.5} \quad \text{s.t.} \quad c_f + c_m = (1+\alpha)[(w_f + w_m - \phi]$$

The first-order conditions give:

$$c_f + v_f - w_f = c_m + v_m - w_m,$$

and plugging this into the budget constraint yields:

$$c_f(1) = w_f + \frac{\alpha}{2}[w_f + w_m - \phi] - \frac{\phi}{2} + \frac{v_m - v_f}{2},$$
$$c_m(1) = w_m + \frac{\alpha}{2}[w_f + w_m - \phi] - \frac{\phi}{2} + \frac{v_f - v_m}{2}.$$

Utilities are then:

$$u_f(1) = w_f + \frac{\alpha}{2}[w_f + w_m - \phi] - \frac{\phi}{2} + \frac{v_f + v_m}{2},$$
$$u_m(1) = w_m + \frac{\alpha}{2}[w_f + w_m - \phi] - \frac{\phi}{2} + \frac{v_f + v_m}{2}.$$

Given the utilities for a given fertility choice, we can now consider whether the couple will choose to have a child. The female partner prefers to have a child if:

$$u_f(1) \geq u_f(0) \quad \Leftrightarrow \quad v_f + v_m \geq \phi(1+\alpha)$$

The same condition applies to the male partner. Consequently, there is no disagreement, i.e. either both partners want to have a child, or both prefer to remain childless. □

Proof of Proposition 2: We once again characterize the outcome by backward induction. In the case without children, the resource allocation in the marriage solves the maximization problem:

$$\max_{c_f,c_m} \quad [c_f - w_f]^{0.5} [c_m - w_m]^{0.5} \quad \text{s.t.} \quad c_f + c_m = (1+\alpha)[w_f + w_m],$$

which is the same as under the full commitment case. Consequently,

$$c_f(0) = w_f + \frac{\alpha}{2}[w_f + w_m] \quad \text{and} \quad c_m(0) = w_m + \frac{\alpha}{2}[w_f + w_m],$$

and utilities are:

$$u_f(0) = w_f + \frac{\alpha}{2}\left[w_f + w_m\right] \quad \text{and} \quad u_m(0) = w_m + \frac{\alpha}{2}\left[w_f + w_m\right].$$

In the case with children, the maximization problem to determine the resource allocation is now different, because bargaining takes place ex post, with the new outside options given the presence of a child:

$$\max_{c_f, c_m} \quad \left[c_f - (w_f - \chi_f\phi)\right]^{0.5}\left[c_m - (w_m - \chi_m\phi)\right]^{0.5} \quad \text{s.t.} \quad c_f + c_m = (1+\alpha)\left[(w_f + w_m - \phi\right].$$

First-order conditions now give us:

$$c_f - (w_f - \chi_f\phi) = c_m - (w_m - \chi_m\phi),$$

and plugging this into the budget constraint yields:

$$c_f(1) = w_f - \chi_f\phi + \frac{\alpha}{2}\left[(w_f + w_m - \phi\right], \tag{23}$$

$$c_m(1) = w_m - \chi_m\phi + \frac{\alpha}{2}\left[(w_f + w_m - \phi\right]. \tag{24}$$

Utilities then are:

$$u_f(1) = w_f - \chi_f\phi + \frac{\alpha}{2}\left[(w_f + w_m - \phi\right] + v_f \quad \text{and}$$
$$u_m(1) = w_m - \chi_m\phi + \frac{\alpha}{2}\left[(w_f + w_m - \phi\right] + v_m.$$

We can now move to the first stage and characterize the fertility preferences of the two spouses. The wife wants to have a child if:

$$u_f(1) \geq u_f(0) \quad \Leftrightarrow v_f \geq \left(\chi_f + \frac{\alpha}{2}\right)\phi,$$

and the male partner would like to have a child if:

$$u_m(1) \geq u_m(0) \quad \Leftrightarrow v_m \geq \left(\chi_m + \frac{\alpha}{2}\right)\phi.$$

In these inequalities, the term $\chi_g\phi$ represents the direct cost of having the child to spouse g. Since bargaining is ex post, having a child lowers the outside option, so that (unlike in the commitment solution) the spouse bearing the greater burden of child care is not compensated. The second term $(\alpha/2)\phi$ represents the loss in marital surplus due to the

cost of a child. This part of the cost of childbearing is shared equally between the spouses.

Depending on v_f and v_m, it is possible that neither, both, or just one of the spouses would like to have a child. Hence, in the case of limited commitment disagreement between the two partners about fertility is possible. □

Proof of Proposition 3: Fertility preferences for gender $g \in \{f, m\}$ have independent uniform density on $\mu_g - (d_g)^{-1}/2, \mu_g + (d_g)^{-1}/2$. The distribution function is given by (in the relevant range):

$$F(v_f, v_m) = \left(v_f - \left(\mu_f - \frac{1}{2d_f}\right)\right) d_f \left(v_m - \left(\mu_m - \frac{1}{2d_m}\right)\right) d_m,$$

and the fraction of couples who have a child is given by:

$$\begin{aligned} E(b) = 1 - \left(\tilde{v}_f - \left(\mu_f - \frac{1}{2d_f}\right)\right) d_f - \left(\tilde{v}_m - \left(\mu_m - \frac{1}{2d_m}\right)\right) d_m \\ + \left(\tilde{v}_f - \left(\mu_f - \frac{1}{2d_f}\right)\right) d_f \left(\tilde{v}_m - \left(\mu_m - \frac{1}{2d_m}\right)\right) d_m. \end{aligned} \quad (25)$$

Given (9) and (10), the average fertility rate is a quadratic and concave function of the female cost share χ_f (i.e., the quadratic term has a negative sign). The derivative of average fertility with respect to χ_f is:

$$\begin{aligned} \frac{\partial E(b)}{\partial \chi_f} = \phi d_m \left[1 - \left((\chi_f + \alpha/2)\,\phi - \left(\mu_f - \frac{1}{2d_f}\right)\right) d_f\right] \\ - \phi d_f \left[1 - \left((1 - \chi_f + \alpha 2)\,\phi - \left(\mu_m - \frac{1}{2d_m}\right)\right) d_m\right], \end{aligned} \quad (26)$$

which simplifies to:

$$\frac{\partial E(b)}{\partial \chi_f} = \phi(d_m - d_f) + \phi d_f d_m \left[(1 - 2\chi_f)\phi + \mu_f - \mu_m + \frac{1}{2}\left(\frac{1}{d_m} - \frac{1}{d_f}\right)\right].$$

Equating the right-hand side to zero gives the cost share $\hat{\chi}_f$ at which fertility is maximized (assuming that the solution is interior):

$$\hat{\chi}_f = \frac{1}{2} + \frac{1}{2\phi}\left[\mu_f - \mu_m + \frac{1}{2}\frac{d_m - d_f}{d_f d_m}\right]. \quad (27)$$

Taking corner solutions into account, the fertility maximizing cost share is given by expression (14) in the statement of the proposition. Moreover, starting with (26), if there is

an interior maximum we have:

$$\phi d_m \left[1 - \left(\left(\hat{\chi}_f + \alpha/2\right)\phi - \left(\mu_f - \frac{1}{2d_f}\right)\right) d_f\right] \\ = \phi d_f \left[1 - \left(\left(1 - \hat{\chi}_f + \alpha 2\right)\phi - \left(\mu_m - \frac{1}{2d_m}\right)\right) d_m\right],$$

and hence:

$$\frac{d_f}{d_m} = \frac{1 - \left(\left(\hat{\chi}_f + \alpha/2\right)\phi - \left(\mu_f - \frac{1}{2d_f}\right)\right) d_f}{1 - \left(\left(1 - \hat{\chi}_f + \alpha 2\right)\phi - \left(\mu_m - \frac{1}{2d_m}\right)\right) d_m} = \frac{1 - F_f(\tilde{v}_f)}{1 - F_m(\tilde{v}_m)}.$$

Thus, as stated in the last part of the proposition, if the distributions of female and male child preferences have different densities, fertility is maximized if the ratio of densities is equal to the fraction of individuals agreeing to have a child for each gender. □

Proof of Proposition 4: The second period of the two-period model is formally identical to the static model analyzed in Proposition 2, and hence conditions (7) and (8) are applicable, which gives (16) and (17). The expected utilities in period 2 as a function of first-period utility are then given by:

$$V_g(b_1) = \int_{v_{f,2}} \int_{v_{m,2}} \Bigg[w_g + \frac{\alpha}{2}\left(w_f + w_m\right) \\ + I(v_{f,2} \geq \tilde{v}_{f,2}, v_{m,2} \geq \tilde{v}_{m,2}) \left(v_{g,2} - \left(\chi_g + \frac{\alpha}{2}\right)\phi\right) \Bigg] f(v_{f,2}, v_{m,2}|b_1)\, dv_{f,2}\, dv_{m,2}, \quad (28)$$

where $f(v_{f,2}, v_{m,2}|b_1)$ is the joint density of fertility preferences in the second period given b_1. Given these utilities, the terms $V_g(1) - V_g(0)$ then represent the change in second period expected utility as a function of the initial fertility choice. From the perspective of deciding on fertility in the first period, these terms act like a constant that adds to (or subtract from) the benefit of children. Applying Proposition 2, the conditions for having a baby in the first period are then:

$$v_{f,1} + \beta\left(V_f(1) - V_f(0)\right) \geq \left(\chi_f + \frac{\alpha}{2}\right)\phi, \\ v_{m,1} + \beta\left(V_m(1) - V_m(0)\right) \geq \left(\chi_m + \frac{\alpha}{2}\right)\phi,$$

which gives (18) and (19). □

Proof of Proposition 5: Given that fertility preferences in the second period do not de-

pend on the fertility realization in the first period, we have $V_f(0) = V_f(1)$ and $V_m(0) = V_m(1)$. Hence, given Proposition 4 the conditions for fertility in each period are the same as those for the single period model characterized in Proposition 2. We therefore obtain the same fertility rate in both periods, $E(b_1) = E(b_2)$, and Proposition 3 applies to each period separately. □

Proof of Proposition 6: We proceed by backward induction. If $b_1 = 1$, we have $v_{f,2} = v_{m,2} = 0$. Given (16) and (17), this guarantees that no additional child will be born in the second period, and second-period utilities are (given Nash bargaining):

$$V_f(1) = w_f + \frac{\alpha}{2}(w_f + w_m),$$
$$V_m(1) = w_m + \frac{\alpha}{2}(w_f + w_m).$$

Conversely, if we have $b_1 = 0$, the preference realizations $v_{f,2} = v_{m,2} = (1 + \alpha/2)\phi$ guarantees that the conditions (16) and (17) are satisfied, so that $b_2 = 1$ for sure. We therefore have $b_2 = 1 - b_1$ and, in expectation:

$$E(b_2) = 1 - E(b_1),$$

which gives (20) and (20). Continuing, the resulting second-period utilities conditional on $b_1 = 0$ are:

$$V_f(0) = w_f - \chi_f\phi + \frac{\alpha}{2}(w_f + w_m - \phi) + \left(1 + \frac{\alpha}{2}\right)\phi,$$
$$V_m(0) = w_m - \chi_m\phi + \frac{\alpha}{2}(w_f + w_m - \phi) + \left(1 + \frac{\alpha}{2}\right)\phi,$$

which can be simplified to:

$$V_f(0) = w_f + (1 - \chi_f)\phi + \frac{\alpha}{2}(w_f + w_m),$$
$$V_m(0) = w_m + (1 - \chi_m)\phi + \frac{\alpha}{2}(w_f + w_m).$$

Given these utilities, the impact of having a child in the first period on continuation utility is:

$$V_f(0) - V_f(1) = (1 - \chi_f)\phi,$$
$$V_m(0) - V_f(1) = \chi_f\phi.$$

We now move to the fertility decision in the first period. The conditions (18) and (19) are:

$$v_{f,1} \geq \left(\chi_f + \frac{\alpha}{2}\right)\phi + \beta(1-\chi_f)\phi,$$
$$v_{m,1} \geq \left(1-\chi_f + \frac{\alpha}{2}\right)\phi + \beta\chi_f\phi,$$

which can be written as:

$$v_{f,1} \geq \left(\beta + (1-\beta)\chi_f + \frac{\alpha}{2}\right)\phi,$$
$$v_{m,1} \geq \left(\beta + (1-\beta)(1-\chi_f) + \frac{\alpha}{2}\right)\phi.$$

With the change of variables:

$$\tilde{\phi} = (1-\beta)\phi,$$
$$\tilde{\alpha} = (\alpha + 2\beta)/(1-\beta),$$

the conditions can be written as:

$$v_{f,1} \geq \left(\chi_f + \frac{\tilde{\alpha}}{2}\right)\tilde{\phi},$$
$$v_{m,1} \geq \left(1-\chi_f + \frac{\tilde{\alpha}}{2}\right)\tilde{\phi}.$$

The conditions therefore are of the form (7) and (8), so that the results in Proposition 3 apply with the transformed parameters $\tilde{\phi}$ and $\tilde{\alpha}$. □

A.2 Correlated Child Preferences

[TO BE WRITTEN]

A.3 Altruism as an Explanation of Link Between Fertility Intentions and Fertility Outcomes

Our baseline model assumes that a child is born only if the spouses agree on having a child. In the data, however, we find that some children are born even when there is disagreement. In the quantitative model in Section 5, we specify fixed probabilities of

childbirth in the case of disagreement. One interpretation of these probabilities is that one spouse may be able to deceive the other in terms of the use of birth control. In this appendix, we develop an alternative interpretation, based on altruism between the spouses.

In the case of altruism, we specify the value function of a married spouse as

$$v_f(n) = u_f(n) + \lambda u_m(n) \quad \text{and} \quad v_m(n) = u_m(n) + \lambda u_f(n),$$

where u_f and u_m denote individual utilities. Consequently for couples with children we have

$$v_f(1) = c_f + v_f + \lambda \left[c_m + v_m\right] \quad \text{and} \quad v_m(1) = c_m + v_m + \lambda \left[c_f + v_f\right].$$

In the case without children, we have

$$v_f(0) = c_f + \lambda c_m \quad \text{and} \quad v_m(0) = c_m + \lambda c_f.$$

We assume that couples are still altruistic towards each other, even when they don't reach an agreement through bargaining. Consequently, outside options are

$$\begin{aligned} v_f^d(1) &= (1-\phi)w_f + v_f + \lambda \left[w_m + v_m\right] \quad \text{and} \\ v_m^d(1) &= w_m + v_m + \lambda[(1-\phi)w_f + v_f] \quad \text{as well as} \\ v_f^d(0) &= w_f + \lambda w_m \quad \text{and} \quad u_m^d(0) = w_m + \lambda w_f. \end{aligned}$$

For simplicity we let the altruism parameter λ be identical for both partners.

Without children In the case without children, the resource allocation of the couple solves the maximization problem

$$\begin{aligned} \max_{c_f, c_m} \quad & \left[c_f + \lambda c_m - w_f - \lambda w_m\right]^{0.5} \left[c_m + \lambda c_f - w_m - \lambda w_f\right]^{0.5} \\ \text{s.t.} \quad & c_f + c_m = (1+\alpha)\left[w_f + w_m\right], \end{aligned}$$

First order conditions again give us

$$c_f - w_f = c_m - w_m$$

and plugging this into the budget constraint yields the known solution to the limited commitment model

$$c_f(0) = w_f + \frac{\alpha}{2}\left[w_f + w_m\right] \quad \text{and} \quad c_m(0) = w_m + \frac{\alpha}{2}\left[w_f + w_m\right].$$

Individual utilities are

$$u_f(0) = w_f + \frac{\alpha}{2}\left[w_f + w_m\right] \quad \text{and} \quad u_m(0) = w_m + \frac{\alpha}{2}\left[w_f + w_m\right],$$

but the value function for each partner is

$$\begin{aligned} v_f(0) &= w_f + \lambda w_m + \frac{(1+\lambda)\alpha}{2}\left[w_f + w_m\right] \quad \text{and} \\ v_m(0) &= w_m + \lambda w_f + \frac{(1+\lambda)\alpha}{2}\left[w_f + w_m\right]. \end{aligned}$$

With children In the case with children the first order condition again reads

$$c_f - (1-\phi)w_f = c_m - w_m$$

and plugging this into the budget constraint yields the limited commitment allocations

$$\begin{aligned} c_f(1) &= (1-\phi)w_f + \frac{\alpha}{2}\left[(1-\phi)w_f + w_m\right] \quad \text{as well as} \\ c_m(1) &= w_m + \frac{\alpha}{2}\left[(1-\phi)w_f + w_m\right]. \end{aligned}$$

Individual utilities then are

$$\begin{aligned} u_f(1) &= (1-\phi)w_f + \frac{\alpha}{2}\left[(1-\phi)w_f + w_m\right] + v_f \quad \text{and} \\ u_f(1) &= w_m + \frac{\alpha}{2}\left[(1-\phi)w_f + w_m\right] + v_m, \end{aligned}$$

but the value function for each partner is

$$\begin{aligned} v_f(1) &= (1-\phi)w_f + v_f + \lambda(w_m + v_m) + \frac{(1+\lambda)\alpha}{2}\left[(1-\phi)w_f + w_m\right] \quad \text{and} \\ v_m(1) &= w_m + v_m + \lambda\left[(1-\phi)w_f + v_f\right] + \frac{(1+\lambda)\alpha}{2}\left[(1-\phi)w_f + w_m\right]. \end{aligned}$$

Decision about children We now say that partners decide to have a children if their value function in this case is larger than when not having a kid. For the female partner

this means

$$v_f(1) \geq v_f(0) \quad \Leftrightarrow \quad v_f + \lambda v_m \geq \left[1 + \frac{\alpha}{2} + \lambda\frac{\alpha}{2}\right] \phi w_f$$

and for the male partner

$$v_m(1) \geq v_m(0) \quad \Leftrightarrow v_m + \lambda v_f \geq \left[\frac{\alpha}{2} + \lambda\left(1 + \frac{\alpha}{2}\right)\right] \phi w_f.$$

Note that the decision for children is equal to the pure limited commitment case if $\lambda = 0$ and equal to the full commitment case if $\lambda = 1$.

Importance of having children In contrast to the decision about children, we say that having a child is important for a partner, if her/his individual utility from the child is larger than in a situation without kids. Therefore a female partner will intend to have a child if

$$u_f(1) \geq u_f(0) \quad \Leftrightarrow \quad v_f \geq \left[1 + \frac{\alpha}{2}\right] \phi w_f.$$

and a male partner if

$$u_m(1) \geq u_m(0) \quad \Leftrightarrow \quad v_m \geq \frac{\alpha}{2} \phi w_f.$$

Note that these inequalities are the same as in the limited commitment model. Consequently, when expressing their intention about children, partners still take into account that they are potentially worse of from having a child than from not having one. Yet, if only one partner wants a child hard enough, the altruism of the other partner might lead to a decision in favor of children, although this other partner might incur an individual utility loss.

References

Adda, Jérôme, Christian Dustmann, and Katrien Stevens. 2015. "The Career Costs of Children." Unpublished Manuscript, Bocconi University.

Ashraf, Nava, Erica Field, and Jean Lee. 2014. "Household Bargaining and Excess Fertility: An Experimental Study in Zambia." *American Economic Review* 104 (7): 2210–2237.

Barro, Robert J., and Gary S. Becker. 1989. "Fertility Choice in a Model of Economic Growth." *Econometrica* 57 (2): 481–501.

Basu, Kaushik. 2006. "Gender and Say: A Model of Household Behavior with Endogenously Determined Balance of Power." *Economic Journal* 116 (511): 558–580.

Baudin, Thomas, David de la Croix, and Paula Gobbi. 2015. "Fertility and Childlessness in the United States." Forthcoming, American Economic Review.

Becker, Gary S. 1960. "An Economic Analysis of Fertility." *Demographic and Economic Change in Developed Countries*. Princeton: Princeton University Press.

Becker, Gary S., and Robert J. Barro. 1988. "A Reformulation of the Economic Theory of Fertility." *Quarterly Journal of Economics* 103 (1): 1–25.

Blundell, Richard, Pierre-André Chiappori, and Costas Meghir. 2005. "Collective Labor Supply with Children." *Journal of Political Economy* 113 (6): 1277–1306.

Caucutt, Elizabeth M., Nezih Guner, and John Knowles. 2002. "Why Do Women Wait? Matching, Wage Inequality, and the Incentives for Fertility Delay." *Review of Economic Dynamics* 5 (4): 815–855.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2012. "Married with Children: A Colective Labor Supply Model with Detailed Time Use and Intrahousehold Expenditure Information." *American Economic Review* 102 (7): 3377–3405.

Chiappori, Pierre-André. 1988. "Rational Household Labor Supply." *Econometrica* 56 (1): 63–90.

———. 1992. "Collective Labor Supply and Welfare." *Journal of Political Economy* 100 (3): 437–467.

Echevarria, Cristina, and Antonio Merlo. 1999. "Gender Differences in Education in a Dynamic Household Bargaining Model." *International Economic Review* 40 (2): 265–286.

Eswaran, Mukesh. 2002. "The Empowerment of Women, Fertility, and Child Mortality: Towards a Theoretical Analysis." *Journal of Population Economics* 15 (2): 433–454.

Feyrer, James, Bruce Sacerdote, and Ariel Dora Stern. 2008. "Will the Stork Return to Europe and Japan? Understanding Fertility within Developed Nations." *The Journal of Economic Perspectives* 22 (3): 3–22.

Fisher, Hayley. 2012. "Fertility in a Collective Household Model." Unpublished Manuscript, University of Sydney.

Galor, Oded, and David N. Weil. 1996. "The Gender Gap, Fertility, and Growth." *American Economic Review* 86 (3): 374–387.

Greenwood, Jeremy, Nezih Guner, and John Knowles. 2003. "More on Mariage, Fertility, and the Distribution of Income." *International Economic Review* 44 (3): 827–862.

Greenwood, Jeremy, Ananth Seshadri, and Guillaume Vandenbroucke. 2005. "The Baby Boom and Baby Bust." *American Economic Review* 95 (1): 183–207.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. "Engines of Liberation." *Review of Economic Studies* 72 (1): 109–133.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2014. "Childcare Subsidies and Household Labor Supply." Unpublished Manuscript, Arizona State University.

Guner, Nezih, and John Knowles. 2009. "Why is the Rate of Single-Parenthood Lower in Canada than in the U.S.? A Dynamic Equilibrium Analysis of Welfare Policies." *Canadian Journal of Economics* 42 (1): 56–89.

Hener, Timo. 2010. "Do Couples Bargain over Fertility? Evidence based on Child Preference Data." SOEPpapers on Multidisciplinary Panel Data Research 323.

Iyigun, Murat, and Randall P. Walsh. 2007. "Endogenous Gender Power, Household Labor Supply and the Demographic Transition." *Journal of Development Economics* 82 (1): 138–155.

Jones, Larry E., and Alice Schoonbroodt. 2010. "Baby Busts and Baby Booms: The Fertility Response to Shocks in Dynastic Models." NBER Working Paper 16596.

Kemnitz, Alexander, and Marcel Thum. 2014. "Gender Power, Fertility, and Family Policy." *Scandinavian Journal of Economics* 117 (1): 220–247.

Klawon, Emily, and Jill Tiefenthaler. 2001. "Bargaining over Family Size: The Determinants of Fertility in Brazil." *Population Research and Policy Review* 20 (5): 423–40.

Lundberg, S., and R. A. Pollak. 1993a. "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy* 101 (6): 988–1010.

Lundberg, S., R. A. Pollak, and T.J. Wales. 1997. "Do Husbands and Wives Pool Their Resource? Evidence from the United Kingdom Child Benefit." *Journal of Human Resources* 32 (3): 463–480.

Lundberg, Shelly, and Robert A. Pollak. 1993b. "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy* 101 (6): 988–1010.

Mazzocco, Maurizio. 2007. "Household Intertemporal Behavior: A Collective Characterization and a Test of Commitment." *Review of Economic Studies* 74 (3): 857–895.

Rasul, Imran. 2008. "Household Bargaining over Fertility: Theory and Evidence from Malaysia." *Journal of Development Economics* 86 (2): 215–241.

Westoff, Charles F. 2010. "Desired Number of Children: 2000-2008." DHS Comparative Reports No. 25.