

Birth Control and Female Empowerment: An Equilibrium Analysis¹

Pierre-André Chiappori² Sonia Oreffice³

February 2006

Please do not quote

¹Paper presented at Boston University, CEMFI, Columbia University, Hunter College, Clemson University, University of British Columbia, University of Victoria, and at the 2005 World Congress of the Econometric Society. We thank the participants and Gary Becker, Guillermo Caruana, Kevin Lang, Claudio Michelacci, Javier Suarez and Derek Neal for many helpful discussions. Financial support from the NSF (award # SES-05-32398) is gratefully acknowledged. Errors are ours.

²E. Rowan and Barbara Steinschneider Professor of Economics, Columbia University. Email: pc2167@columbia.edu

³Department of Economics, Clemson University. Email: oreffic@clemson.edu

Abstract

It has been widely argued that innovations in birth control technology during the last decades have affected not only women's fertility choices, but generally their position in families and society. We analyze, from a theoretical perspective, the impact of these innovations on the intrahousehold allocation of resources. We consider a model of frictionless matching on the marriage market in which women differ in their preferences for children. The qualitative features of the equilibrium depend on whether men or women are in excess supply on the marriage market. If women are in excess supply, then an improvement in the birth control technology generally increases the 'power', hence the welfare, of all women, including those who are not interested in the new technology. However, this 'empowerment' effect crucially relies on the availability of the new technology to single women. Should the innovation be reserved to married women (as was initially the case for the pill), the conclusion may be reversed, resulting in a 'disempowerment' effect. Finally, various extensions are considered.

1 Introduction

The innovations in birth control during the last four decades have had a strong impact on modern societies. The direct effect on demography, and especially on birth rates in specific social groups, has been abundantly studied.¹ An additional, indirect effect operates through the reduction in uncertainty faced by couples and single women due to the increased ability to plan demographic phenomena. This fact has dramatically altered the context in which decisions regarding human capital accumulation were made. For instance, many authors have argued that the spectacular increase in women's education and participation to the labor market since the 60s was largely due to these innovations.²

A third, and relatively less studied consequence goes through a different channel - namely, the impact of the changes on the 'balance of power' within the couple. After all, innovations in birth control technology, including the legalization of abortion, have a potentially huge effect on men's and women's respective decision rights within the household regarding such crucial issues

¹See for instance Phillip B. Levine et al. (1999), Jacob Alex Klerman (1999), Joshua Angrist and Evans (1996) and Elizabeth Oltmans Ananat, Jonathan Gruber and Phillip B. Levine (2004).

²See Robert Michael (2000) and Claudia Goldin and Lawrence Katz (2002).

as the number or timing of births. It is hard to believe that a shift of such magnitude would leave the balance of powers unaltered. This claim has been put forth by a number of sociologists³, but there is little in terms of an economic analysis of this phenomenon.

Sonia Oreffice (2005) has provided an econometric study based on the collective approach to household behavior. The basic idea is to check whether legalizing abortion affected the labor supply of married females in their fertile age. If, as is often argued, the legalization actually improved the wife's bargaining situation within married couples, then she should receive a larger share of household resources. By a standard income effect, the reform should thus decrease her labor supply and increase her husband's. As it turns out, the predictions of the model are remarkably well confirmed by the data. A significant effect is observed on both male and female labor supply for married couples, but not for older couples nor for single males and females.

Although the claim that abortion legalization influenced intrahousehold bargaining has received some empirical support, the mechanism by which this empowerment occurs still deserves some scrutiny. While it is not hard to convince oneself that *some* women will gain, whether *all* will is another

³See for instance Williams (1994), H eritier (2002), or Coombs and Fernandez (1978).

matter. A strong objection is that women have heterogeneous preferences for fertility (or different attitudes toward abortion); some, in particular, do not consider abortion as an option, either for religious and ethical reasons or because they do want children. Whether the legalization will benefit these women as well is not clear. From a general perspective, the new context will affect the matching process on the market for marriage, and in particular the way the surplus generated by marriage is shared between members. In principle, such ‘general equilibrium’ effects could annihilate or even reverse the direct impact of the reform, particularly for these women who are unlikely to derive much direct benefit from it.

The goal of the present paper is precisely to construct an explicit model of matching on the marriage market in which these phenomena can be formally analyzed. We use a barebone setting in which, in the initial situation, any woman who wants a child is biologically able to have one, but unwanted pregnancies are possible. Children (whether wanted or not) generate utility for both parents, but decrease the mother’s ability to earn income. We assume that men are identical, while women differ by their preferences towards children. Finally, marriage is modelled as the outcome of a frictionless matching process; utilities are assumed transferable, so that the impact of the matching

process on individual utilities can readily be assessed. Our key assumption is that households are able to commit, at least partially, on the division of resources during marriage. It follows that the the environment at the date of marriage, which determines the agreement reached between spouses when they marry, has potentially lasting effects on intrahousehold allocation. This Beckerian assumption has in our context crucial consequences. As it is common in this type of frameworks, the property of the equilibrium depends on the respective proportions of male and females on the marriage market. A more specific property of our setting is the existence of three classes of women: those who like children enough to be willing to be a mother even when single, those who want to have a child only if married, and those who never want to have a child.⁴ In particular, when women are globally in excess supply, then the main qualitative patterns of the stable match depend on which class the 'marginal' married woman belongs to. it is important to note that the definition of these classes is *endogenous*, and may be affected by various external factors, including social policies.

The framework generates interesting comparative statics results. A pro-

⁴Our simple, *static* framework cannot address such dynamic issues as the timing of birth. For instance, deciding not to have a child can sometimes be interpreted as postponing the time of birth.

portional increase in female income (both with and without children) decreases fertility; moreover, the gain for the wife typically exceed the additional income because of the intrahousehold redistribution triggered by the change. A decline in the supply of men on the marriage market decreases total fertility but tends to increase out-of-wedlock births; it always harms women's well-being. The introduction of more generous benefits targeted to single mothers also increases out-of-wedlock fertility, but the impact on female welfare is more complex, since its consequences on the qualitative patterns of the stable match (and in particular on intra-household allocation of the surplus generated by marriage) must be taken into account. We argue that these features are broadly compatible with the main changes that took place in the US over the last decades.

We then consider the introduction of a new technology that allows women (or couples) to better control fertility at no cost. Note that our general setting can be applied not only to abortion, but to any technological innovation affecting birth control. A particularly interesting example is provided by the pill, which first became available to married women only, and later to singles as well. Our main finding is that in our simplified setting, the legalization of abortion (and generally any improvement in birth control that is freely

available to *all* women) never worsens women's well being (and never increase men's). If, in particular, the marriage market is characterized by an excess supply of women that is not too 'large' (in a sense precisely defined below), all women are strictly better off after the reform. Even women who do not actually use the new technology benefit from its introduction, because they are able to grab a larger share of household resources; in some cases, these women are actually those who gain the most. Correspondingly, men are strict losers from the reform. Lastly, these conclusions remain essentially valid when individuals' access to the technology has a cost; these aspects are precisely discussed.

Our model thus predicts that a better birth control technology leads to an empowerment of all women in the economy. However, this 'empowerment effect' relies on a market adjustment mechanism in which single women play a key role. Indeed, the reform benefits married women with children *only insofar as the new technology is available to single women as well*. Paradoxically, should the new technology be exclusively reserved to married women (as was initially the case for the pill), then for most (married) women the impact of the reform in terms of intrahousehold allocation is either nil or even negative, at least when women are in excess supply.

Related literature Our work is related to two lines of research. From a theoretical perspective, modeling marriage markets as matching equilibria is a standard approach, used in several contributions.⁵ Related, although different, approaches rely on search models, in which initial matching is generally random, but agents may decide to accept the current match or continue searching;⁶ matching equilibria can be seen as limit cases of search models when frictions become negligible (so that repeated sampling is costless). An important aspect is that both approaches share the property that the environment at the date of marriage is a driving factor in the determination of intrahousehold allocation; as such, they both require some level of commitment.

An alternative view is based on intrahousehold bargaining, and postulates that household decision is the outcome of a bargaining game in which commitment is limited or absent; it follows that any change in the environment will be reflected in the household's decision, irrespective of the agreement at

⁵Following Gary Becker's (1973, 1974, 1993) seminal work, one can mention, among others, Dale Mortensen (1988), Alvin E. Roth and Marilda Oliveira Sotomayor (1990), Murat Iyigun and Randy Walsh (2004), Eugene Choo and Aloysius Siow (2005), and Pierre-André Chiappori, Murat Iyigun and Yoram Weiss (2006).

⁶Examples include Kenneth Burdett and Melvyn G. Coles (1997, 1999), S. Rao Aiyagari, Jeremy Greenwood and Nezih Guner (2000), Robert Shimer and Lones Smith (2000) and Chiappori and Weiss (2003, 2005).

the date of marriage.⁷ As we shall see, our results are quite different from those stemming from a bargaining approach.

Our work is also related to the contribution of Akerlof, Yellen and Katz (1996, from now on AYK) on out-of-wedlock child bearing in the United States. The relationship between their approach and ours is discussed in more detail at the end of this paper. At this stage, we may just emphasize a crucial difference. In AYK, the level of welfare reached by *individuals* within each couple is taken as exogenously given, whereas it is fully endogenous in our context. Our goal is precisely to analyze how equilibrium conditions on the marriage market influence intrahousehold allocation of welfare. In that sense, our approach is closer to that proposed by Neal (2004); in particular, Neal also constructs a model in which intrahousehold allocation of resources is the outcome of a matching equilibrium. The major difference with our paper is that Neal emphasizes income heterogeneity, whereas differences in taste for children are the driving force in our framework. Hence our approach is largely complementary to those existing in the literature, although some conclusions differ.

The paper is organized as follows: Section 2 describes the model; Section

⁷See for instance Shelly Lundberg and Robert A. Pollak (1996) for a survey.

3 provides the equilibrium analysis and some comparative statics; Section 4 studies the effects of the technological innovation on the equilibrium outputs; Section 5 presents some extensions and discusses the links with the existing literature; Section 6 concludes the paper.

2 The model

Preferences and budget constraints In the economy we consider, there exists a continuum of men and women who derive utility from one private composite good a and from children. For the sake of expositional simplicity, we only consider the choice between having children or not, although the generalization to different numbers of children is straightforward. Let the dummy variable k denote the presence ($k = 1$) or the absence ($k = 0$) of children in the household; alternatively, k can be interpreted as the additional child that men and women may have at any given moment.

Men have identical, quasi-linear preferences over consumption and children. To sharpen our analysis, we assume that the utility of single men only depends on their consumption; i.e., men cannot derive utility from (and do not share the costs of) out-of-wedlock children, due to the fact that they do

not live in the same household.⁸ On the other hand, married men’s utility is of the form $U_H(a_H, k) = a_H + u_H \cdot k$, where the parameter $u_H > 0$ is identical for all men in the economy. Note, in particular, that the utility of a married man does not depend on the identity of the person he marries, but only on the joint fertility decision and on the share of composite good he receives.⁹

Women differ in their preferences toward children. Specifically, female utility functions take the quasi-linear form $U(a, k) = a + uk$. Here, each woman is characterized by the individual-specific taste parameter u , which is distributed according to the density f and the CDF F over some support included in the interval $[0, U]$.¹⁰ Note, in particular, that utilities are transferable, a property that will facilitate the matching analysis that follows. We assume that any woman (single or married) who wants a child can have one. However, should she plan not to have children, unwanted births may still

⁸We can modify these assumptions to allow for single men sharing some costs of children (in terms of child support) and deriving some utility from children who are not living with them; the qualitative conclusions remain unchanged.

⁹Following Chiappori and Weiss (2002), one could include an additional, match-specific non-monetary benefit θ_i received by both spouses from the companionship in marriage. Results would not change provided that the distribution of θ_i is the same for women of different types.

¹⁰We assume that taste parameters are publicly observable; say, a man, when dating a woman, acquires a reasonably good understanding of her attitude toward children. Also, note that we do **not** make assumptions on the respective values of males’ and females’ preferences for children (i.e., technically, on the location of u_H with respect to the support of the distribution of u). Our model allows for all women caring more for children than any man, or any alternative assumption.

occur with some probability p , which depends on the technology available.¹¹

Throughout the paper, the price of the private composite good is normalized to 1. Male income is denoted Y . Women initially receive an income y ; however, if a woman has children, her income drops to z , with $z < y$, reflecting both the loss in her earning capacity due to childbearing and the cost of raising the child. Hence a single woman without children consumes her income y ; if she decides to have a child (or if an unwanted pregnancy occurs), she also consumes her income (which has dropped to z) and receives moreover the utility u from her child, which is independent of her marital status.

Regarding couples, we assume that $u_H < y - z$, i.e. that the gain received by the husband from having a child does not offset by itself the loss in income experienced by the wife. This assumption implies, in our framework, that the couple's decision to have a child or not will also depend on the wife's preferences. Therefore married women must agree with their husband on two issues. One is the fertility decision; i.e., they must decide whether

¹¹Introducing uncertainty about the ability of having children would not change our qualitative results. Indeed, women in general do not know whether they are biologically able to have children until they try to have one; hence, there would be no hidden information or selection into marriage in this respect, and in our risk neutral world all payoffs can always be interpreted as expected values. Also, our results are robust to the introduction of different unwanted birth probabilities between married and single women.

to have kids or not, and the decision depends (in particular) on the wife's preferences towards children. The other decision relates to the distribution of resources within the household (i.e., the allocation of total income between male and female consumption of the composite good).¹² Both decision will be ultimately determined by the equilibrium on the market for marriage.

As indicated above, our main focus is the impact of innovations in the birth control technology. Specifically, we model the legalization of abortion (and generally the availability of some birth control technology) as an exogenous decrease in the probability p of experiencing an unwanted pregnancy.¹³

Marriage market We assume that the marriage market is frictionless; i.e., we model marriage as a matching process, for which we will characterize the 'stable' matches in the usual sense (no married person would rather be single, and there cannot be a man and a woman who would both prefer being married together than their current situation). Any individual gets married as long as the utility (s)he can get from marriage is larger than or equal to

¹²Note that should the couple decide not to have children, each member's consumption may also be state-contingent, i.e. depend on whether a child is (unvoluntarily) born or not. In our world, individuals are risk-neutral, and the payoffs can always be interpreted as expected values.

¹³In particular, we disregard for the moment the other benefits derived from the improved birth control technology (e.g., increased ability to plan and control human capital accumulation, or easier access to sexual activity before marriage).

the utility (s)he gets from remaining unmarried. Note that our analysis is consistent with a collective model of household behavior in which spouses interact and take Pareto-efficient decisions, possibly transferring resources to each other. From this perspective, a contribution of this paper, following several others¹⁴, is to endogeneize the 'sharing rule' that characterizes intrahousehold allocation.

Because children are public goods, marriage generates a net surplus; hence it must be the case that either all men or all women are married. As always with matching models, the qualitative properties of the stable match will crucially depend on which side is in excess supply. Throughout the paper, we normalize the mass of women to be 1, and we denote by M the total mass of males on the market. Hence we shall distinguish two main cases, depending on whether the male population M is smaller or larger than 1, as well as several subcases.

¹⁴See for instance Browning, Chiappori and Weiss (2005), Iyigun and Walsh (2004), and Chiappori, Iyigun and Weiss (2005).

3 Stable matching on the marriage market

We now characterize the main features of the equilibrium (or stable assignment) reached on the marriage market.

3.1 Fertility decisions

We first consider the fertility decisions of singles and couples, starting with single individuals. Single men do not make decisions: they consume their income, and get a utility equal to Y . Single women, on the other hand, will decide to have children if and only if the benefit compensates the income loss, i.e. if $u \geq y - z$, leading to a utility equal to $z + u$. In the alternative case when $u < y - z$, single women of type u choose not to have a child; any pregnancy will be involuntary and occur with probability p . Their expected utility will thus be equal to $y(1-p) + p(z + u)$. In what follows, the threshold $y - z$ is denoted \bar{u} ; women whose parameter is larger than or equal to \bar{u} will be said to be of 'high' type.

Regarding couples, note, first, that in a transferable utility context the stable match must maximize total surplus. The total benefit, for a couple, of having a child is $u_H + u$, whereas the cost is $y - z$. It follows that a married couple will plan to have a child if $u \geq y - z - u_H$ - then total utility is

$Y + z + u_H + u$. The threshold $y - z - u_H$ is denoted \underline{u} ; note that $\underline{u} < \bar{u}$. If $u < y - z - u_H$, only unwanted kids are born, leading to an expected total utility equal to $Y + (1 - p)y + p(z + u_H + u)$. Women whose parameter is smaller than \underline{u} will be said to be of 'low' type, while those between \underline{u} and \bar{u} will be called 'intermediate'. To summarize:

- women of 'high' type ($u \geq \bar{u}$) always choose to have a child
- women of 'intermediate' type ($\underline{u} < u < \bar{u}$) choose to have a child only when married
- women of 'low' type ($u \leq \underline{u}$) never choose to have a child

The next step is to analyze the properties of the equilibrium. As we said above, a crucial distinction is related to the respective number of men and women on this market. We thus consider two cases.

3.2 Marriage market with excess supply of women

We first consider the case of an excess supply of women; i.e., we assume that $M < 1$. Because women are in excess supply, they have to compete for marriage. An intuition of the driving force behind this competition is provided by Figure 1, which plots the maximum utility $\Phi(u)$ a man can

achieve when marrying a woman of taste u (in other words, $\Phi(u)$ denotes his utility if he was to appropriate *all* the surplus generated by marriage). The function Φ is increasing; i.e., it is always better (for the husband) to marry a wife with a larger taste coefficient u .

Insert Figure 1 here

More precisely, women whose parameter u is greater than \bar{u} (the 'high' type), and who would plan to have a child even when single, are the most 'attractive' from the male's perspective. While they differ in taste, this difference is irrelevant from a husband's viewpoint, since they require the same compensation c_H for getting married (namely, to be left with a private consumption equal to their income with a child, z). Women between \underline{u} and \bar{u} (the 'intermediate' type) come next in males' preferences. They plan to have a child only when married, and the minimum compensation they require is equal to $c_I(u) = (y - u)(1 - p) + pz$. Note that this compensation decreases with the taste parameter; hence men strictly prefer intermediate women with a higher parameter u . Finally, women with a u smaller than \underline{u} (the 'low' type) never plan to have a child. Again, these women are equivalent

from a husband's perspective, since they require the same compensation for getting married, namely their consumption as single, i.e. $c_L = (1 - p)y + pz$.

As often in matching models, the properties of the stable match crucially depend on the identity of the marginal spouse (i.e., the 'last' married woman). We denote by $u(M)$ the taste parameter of the marginal women (i.e., either the 'last' single woman or the 'first' married woman). Technically, $u(M)$ is defined by the fact that the measure of the set of women with a taste parameter larger than $u(M)$ is exactly M ; i.e., the value $u(M)$ solves the equation

$$\int_{u(M)}^U f(t) dt = M.$$

or equivalently

$$u(M) = F^{-1}(1 - M)$$

Competition between women on the marriage market implies that women who generate a larger surplus for their husband are married in priority. In particular, whenever a woman belonging to the intermediate type is married, then all women with a larger taste parameter are married as well - this is the case depicted in Figure 1a. The intuition is that women with a larger preference for children have a comparative advantage: the compensation they

need from their husband to accept marriage is smaller, because they value highly the prospect of having a child. In general, the category of this marginal woman depends on the location of $u(M)$ with respect to the two thresholds \underline{u} and \bar{u} .

An obvious property of stable matches in this context is that *all males receive the same utility*; indeed, they are assumed identical, and the absence of friction implies that any difference of welfare between males would be competed away.¹⁵ Since the marginal woman is indifferent between being married or single, her husband gets all the surplus generated by the relationship, namely $\Phi(u(M))$. Then all other men receive the same utility. Graphically, this corresponds, on Figure 1, to the horizontal line going through $\Phi(u(M))$.

A crucial insight, at this point, is the following. Take any woman with a taste parameter u larger than $u(M)$. Then *the difference $\Phi(u) - \Phi(u(M))$ represents the surplus received by this woman*.¹⁶ In Figure 1, for instance, the surplus received by any woman of 'high' type is depicted by a bold arrow.

Using this geometric intuition, the characterization of the equilibrium is

¹⁵The case of heterogeneous male population is briefly discussed in the concluding section.

¹⁶If her husband's utility was $\Phi(u)$ he would get all the surplus generated by the marriage. Since his equilibrium utility is only $\Phi(u(M))$, the difference $\Phi(u) - \Phi(u(M))$ represents the part of the surplus appropriated by the wife.

straightforward. Three cases should be distinguished:

- If $M \leq \underline{W} = \int_{\bar{u}}^U f(t) dt$, the excess supply of women is 'large', in the sense that there are less men than high type women (see Figure 1b). Then $u(M) \geq \bar{u}$, and the marginal married woman belongs to the high type. Only (some of) these women are matched. Women of the same type who remain single decide to have a kid; all other women remain single and decide not to have children (although they may have one involuntarily). Regarding welfare issues, note that, in that case, married women receive no surplus from marriage; their consumption is the same as if single.
- If $\underline{W} < M < \bar{W} = \int_{\underline{u}}^U f(t) dt$ ('intermediate' excess supply of women, as depicted in Figure 1a), the marginal wife belongs to the intermediate type. All married women have a child, and consume the same amount, which is such that the marginal wife is indifferent between getting married and remaining single. All married women (but the marginal one) get a positive surplus from marriage, and high type women receive the maximal surplus.
- Finally, when the excess supply of women is small enough (technically,

$M \geq \bar{W}$), the marginal wife belongs to the low type (i.e. $u(M) \leq \underline{u}$ - see Figure 1c). Her fertility is the same with and without marriage - namely, no planned child. Stability requires that her consumption be also the same, i.e. equal to $(1-p)y + pz$. The same conclusion applies to all married, low type women. Other married women belong to the high or intermediate type, hence decide to have a child; their consumption is defined by the fact that men, who are in short supply, must be indifferent between the various potential spouses. Again, this condition generates a positive surplus for all women of high and intermediate types; high type women receive the largest surplus.

These conclusions are formally expressed in the following Proposition:

Proposition 1 *In the presence of excess supply of women, there exists at least one stable match. For any stable match, all men marry; moreover:*

- *if $M \leq \underline{W}$ (Large Excess Supply of Women or LESW), then all married women have a taste parameter $u \geq \bar{u}$; they all choose to have a child, and consume an amount $c_H = z$, thus reaching a utility level of $z + u$. Single women of the high type decide to have a child, while intermediate and low types do not.*

- if $\underline{W} < M < \bar{W}$ (*Intermediate Excess Supply of Women or IESW*), then all married women have a taste parameter $u \geq \underline{u}(M)$, where $\underline{u} < \underline{u}(M) < \bar{u}$. All married women have a child, and consume an amount

$$c_I(M) = (y - u(M))(1 - p) + pz,$$

thus reaching a utility level of

$$(y - u(M))(1 - p) + pz + u.$$

Single women belong to the intermediate and low types; none of them plans to have a child.

- if $M \geq \bar{W}$ (*Small Excess Supply of Women or SESW*), then a married woman decides to have a child if she belongs to the high or intermediate types (i.e., if her taste parameter u is larger than \underline{u}); then she consumes $\bar{c}_C = z + (1 - p)u_H$. If, alternatively, $u \leq \underline{u}$ the (low type) woman does not plan to have a child, and consume an amount $\bar{c}_L = (1 - p)y + pz$. Single women are low type and do not plan to have a child.

The results are summarized in Table 1.

Insert Table 1 here

The equilibrium is technically not unique, because in the first and third situations the identity of married women is indeterminate. In the first case, high type women are identical from the husband's perspective, hence any subset of size M may be married at equilibrium; and the same argument applies to low type women in the third case. Note, however, that the respective consumptions are uniquely determined in each case; so are welfare, since in the first (resp. third) case, high (resp. low) type women are indifferent between marriage and singlehood.

The variation in women's utility across the three types of equilibria exhibits interesting patterns. Not surprisingly, women are better off the smaller their excess supply on the market. However, when women's excess supply is either large or small, their welfare does not depend on the size of the imbalance. In the intermediate case, on the contrary, a marginal increase in the number of men continuously reduces the taste parameter $u(M)$ of the marginal woman, which ameliorates the welfare of all married women.

3.3 Marriage market with excess supply of men

The alternative case in which men are in excess supply on the marriage market is much simpler. All women are married; and stability requires that each married man be indifferent with remaining single. We thus get the following result:

Proposition 2 *For any stable match, all women are married, men are indifferent between being married or single, and their utility is Y . Moreover:*

- *a woman with taste $u \leq \underline{u}$ plans not to have a child; she consumes $(1 - p)y + p(z + u_H)$*
- *a woman with taste $u > \underline{u}$ will have a child and consume $z + u_H$*

In words: since men have to be indifferent between marriage and singlehood, women leave them with whatever is needed for them to achieve their utility as single, namely Y . In particular, the utility men derive from having children is captured by women under the form of additional consumption.

3.4 Comparative statics

Simple as it may be, our model still offers interesting insights on the role of several key parameters.

Incomes Consider, first, the impact of female income on fertility and allocations. Not surprisingly, female utility always increases with both y (her income without a child) and z (her income with a child). However, the respective impact of the two incomes on a woman's welfare depends on the woman's type. For high type women, for instance, an increase in z is always favorable, since they always have a child. Still, y may also matter for them; for instance, in the case of intermediate excess supply of women, the rent received by married women is positively related to the income of childless women. A similar logic prevails for women of the low type: y always matters (they will never *elect* to have a child), but so does z (because of unwanted pregnancies).

The situation is more complex when y and z do not increase by the same *amount*, because of the possible impact on the fertility decisions and the type of equilibrium finally reached. An interesting illustration obtains when both y and z are increased by the same *factor*. Then the difference $y - z$

increases proportionally, which reduces the number of high type women and inflates the low type population (the effect on the intermediate type depends on the distribution of tastes). On Figure 2a, one can see that income growth shifts the function Φ to the right. A first consequence is a negative impact on fertility. More interestingly, the type of equilibrium reached may switch, at least when women are in excess supply; indeed, the thresholds \underline{u} and \bar{u} are positively related to the difference $y - z$. For instance, a 'large' excess supply of women (in the sense just define) may become 'intermediate'; and an intermediate case may switch to 'small'.

We conclude that in a context of excess supply of women, a growth in female income reduces the severity of the excess supply phenomenon, which inflates their welfare gain beyond the change in income, and decreases male well being accordingly. On the contrary, a variation in male income has no impact on fertility and on the equilibrium structure; only men's equilibrium payoffs will vary.

Insert Figure 2 here

Smaller male population Let us now study the impact of the size M of the male population. Not surprisingly, a variation large enough to switch from an excess supply of men to an excess supply of women will favor men and cost to women. More interesting are changes taking place within the excess supply of women regime. Assume, for instance, that the initial situation is characterized by an intermediate excess supply of women (as depicted in Figure 2b). All married women (but the marginal one) receive a share of the surplus generated by marriage. A reduction in the male population, assuming that the equilibrium remains an 'intermediate excess supply' one, moves the threshold $u(M)$ to the right. Some women who were previously married are now single; they decide not to have children, which reduces total fertility, and they lose entirely the share of the marital surplus they initially received. Even women who are married in both cases receive a smaller share of the surplus in the second situation.

Assume, now, that the number of men is reduced below the threshold \underline{W} , so that the equilibrium switches to a 'large excess' type (Figure 2c). Various consequences can be expected:

- In the initial situation, some of the intermediate type women were married and had a child, while they are single and childless (except for

unwanted pregnancies) in the new equilibrium; hence *total fertility is smaller*.

- High type women were all married and had a child in the initial equilibrium. Some of them are now single, but they still choose to have a child. The result is *higher out-of-wedlock fertility*.
- The welfare of *all* married women is decreased down to their reservation utility. Indeed, in the initial, ‘intermediate excess supply’ equilibrium, married women receive a share of the surplus generated by marriage. In the new, ‘large excess supply’ equilibrium, however, men appropriate all the surplus, leaving married women at their reservation utility.
- The same analysis applies to those intermediate women who are married in the initial equilibrium but not in the new one. Again, all but the marginal one receive a fraction of the surplus when married, therefore singlehood results in a welfare loss.
- In both cases, male utility is increased.

Finally, a similar analysis applies if the equilibrium switches from ‘small’ to ‘intermediate excess supply’ types. On the contrary, if the initial equilibrium was characterized by a large excess supply of women, or if their excess

supply is 'small' in both cases, welfare does not change in response to a reduction in M .

Single parent benefits So far, we have assumed that the income of a woman with a child is z whatever her marital status. We now allow for single parent benefits, say b , and assume for simplicity that all single women are eligible. The income of a woman with a child is thus still z when she is married, but becomes $b + z$ when single. We assume, however, that $b \leq u_H$, i.e., that the surplus stemming from the public nature of children in case of marriage compensates the (monetary) loss of the single parent benefit. This condition is necessary in our simplified model for marriage to take place at equilibrium.

The case of an excess supply of men is of little interest: single parent benefits do not matter in that case since all women are married and receive all the surplus created by marriage. We thus consider the alternative context of an excess supply of women.

The main changes with respect to the initial situation can be summarized as follows. First, the upper threshold \bar{u} declines by b : more women are now willing to have children even when single. It follows that when the

excess supply of women is 'large', *out-of-wedlock fertility* increases. Note that, technically, an equilibrium initially classified as 'intermediate' may now become 'large', since the number of high type women has increased.

If we assume that the type of equilibrium does *not* switch, then women always gain from the introduction of the benefits. Specifically:

- In the case of a large excess supply of women, women of intermediate or low type are single and choose not to have a child. Still, they are entitled to the additional payment b in case of unwanted pregnancy; hence an increase in expected utility equal to pb . Women of the high type receive an additional gain equal to b when they are single, because they all have a child. But then equilibrium requires that they receive the same gain when married (their reservation utility has increased by b). We thus have the counter-intuitive result that *the introduction of single-parent benefits profits all women*, and that married women (who do not receive the benefit) gain on average more (per capita) than most singles. Not surprisingly, men lose the same amount b .
- When the excess supply of women is either intermediate or small, all women (married or single) get an extra pb , and all men lose the same expected amount.

These conclusions should however be qualified for two reasons. First, when the initial excess supply of women is 'large', the number of potential spouses (i.e., women of the high type) increases. For a given supply of men, this decreases the probability that each one gets married; note, however, that since women are indifferent between getting married and remaining single, the resulting welfare loss is nil (at least in our model).

Secondly, a more complex situation arises when the change generates a change of equilibrium, the excess supply of women switching either from 'small' to 'intermediate' or from 'intermediate' to 'large'. Indeed, such a switch harms all women, because tougher competition on the marriage market results in smaller rents. The welfare comparison becomes somewhat intricate, because this loss should be traded off with the higher expected utility when single (which boosts married women's reservation point); intuitively, women perceive a smaller rent beyond a higher reservation utility. It can be showed that depending on the parameters, the introduction of the benefit may actually harm some married women.

Finally, if the distribution of these single parent benefits is 'too' generous, in the sense that $b > u_H$, then marriage becomes inefficient from the couple's

viewpoint;¹⁷ indeed, the surplus generated may be inferior to the opportunity cost in terms of lost income. Then welfare benefits would decrease (in our model, eliminate) marriage, fertility becoming exclusively an out-of-wedlock phenomenon.

The US example In practice, our simple model generates predictions that seem to fit fairly well the main stylized facts characterizing the marriage market in the US. Some predictions are fairly straightforward; for instance, the raise of female wages has had a negative impact on fertility. More interesting (and perhaps less expected) is the fact that a decline in the supply of 'marriageable' males reduces total fertility but increases out-of-wedlock birth rate. Such a mechanism may have played an important role in specific sub-markets. For instance, an overwhelming phenomenon that took place during the last decades is the dramatic raise of incarceration rates for young, low-skilled black males. According to Western and Pettit (2000), the percentage of young, high-school dropouts black males in prison or jail reached 36% in 1996, up from 15% in 1982; as a comparison, the rate for young, high-school

¹⁷We assume that the single parent condition can be enforced; i.e., we do not consider situations in which a man and a woman enjoy the benefits of parenthood through cohabitation while receiving the single parent benefit.

dropouts *white* males went only from 4% to 7% over the same period.¹⁸ If we accept that the marriage market is largely assortative by race and education, the 'submarket' of black, low-skilled individuals has experienced a dramatic decline in male supply. According to our analysis, such a brutal change should induce a sharp rise in non marital fertility. Indeed, while the overall fertility rate of black, high-school dropout women aged 25 to 29 has slightly declined over the period (from 102.5 per mil in 1980 to 99.5 per mil in 2000¹⁹), the fraction of women never married with children has raised from 3% in 1960 to 35% in 1990, as compared to 6% for white high-school dropouts of the same age and for black women with college education.²⁰ Our results, on this point, are in line with earlier intuitions proposed by Neal (2004). Also, an implication of our model is that these evolutions may have significantly *decreased* the welfare of unskilled black women over the period.

¹⁸Similar effects can be observed on employment data. Among black high-school dropouts aged 26-36, the percentage of individuals who worked at least 26 weeks during the poll year dropped from 80% in 1960 to 44% in 1990. For high school graduates, the drop is from 86% to 67% (Neal, 2004, table 2). See also Wilson (1987).

¹⁹Computation made by the authors using National Vital Statistics Report, 50-5, 2002, Table 4, p.30; Monthly Vital Statistics Report, 31-8, 1982, Table 19, p. 30; National Vital Statistics Report, 50-5, 2002, Table 21, p. 51; "Educational Attainment, US Census Bureau, Historical Tables, Table A-2, p. A-4.

²⁰The effect may initially have been further boosted by the increased generosity of single parent benefits over the 60s and the early 70s. Note, however, that this latter trend has been largely reverted since then (Moffit 1992), while the increase in birth outside marriage has persisted, suggesting that the marriage market factor may have been dominant.

4 Changes in the birth control technology

We now come to the main implications of our model, namely the impact of a technological change in birth control.

Legalizing abortion We first consider the impact of innovations in birth control technology that reduce the probability of unwanted pregnancies, assuming that all women (including single) are given free access to the technology; a natural example could be the legalization of abortion that took place in the 70s.

We start with the case of an excess supply of men. Then the innovation has no impact on behavior, and a straightforward impact on welfare. Before and after the change, all women are married. The welfare of women of the high and intermediate type does not change: since all are married and planning to have kids, the technology has no consequence for them. Couples in which the wife is of the low type are better off, but only because they do not plan to have a child, and the new technology improves their ability to reach that goal. Clearly, all the gain goes to the wife.

Insert Figure 3 here

The situation is drastically different in the alternative (and probably empirically more relevant) situation in which women are in excess supply. The situation is depicted in Figure 3 (which, for expositional convenience, considers the case in which the risk of unwanted pregnancies goes to zero). The new technology decreases the maximum utility attainable by husbands of low or intermediate type women, resulting in a shifts downwards of the graph of the function Φ . This leads to the following conclusions:

- Not surprisingly, women who do not want to have a child (either because they belong to the low type or because they are single) benefit from the technology, precisely because unwanted pregnancies become less likely. In the extreme situation in which unwanted pregnancies are eliminated, the monetary gain is thus $p(y - z - u)$
- More interesting is the fact that when the excess supply of women is not 'large' (in the sense defined above), women who decide to have a child do also benefit from the technology, although to a lesser extent than singles. The intuition is that the intrahousehold distribution of resources is driven by the marginal women; for a small or intermediate excess supply of women, the marginal woman is indifferent between getting married and remaining single *without kid*. Her reservation util-

ity is thus improved by the new technology. The nature of a matching game, however, implies that any improvement of the marginal agent's situation must be transmitted to all agents 'above' the marginal one.

In practice, in the case of an intermediate excess supply depicted in Figure 3, the benefit experienced by all married women, assuming the new technologies drives the risk of unwanted pregnancies to zero, is $p(y - z - u(M))$ (where, again, $u(M)$ denotes the taste parameter of the marginal married woman). This benefit continuously increases with the number of men M . When the excess supply is small, the gain is pu_H , still smaller than $p(y - z)$ (the gain for single women) but nevertheless positive.

- On the other hand, when the excess supply of women is 'large', married women do *not* benefit from the new technology, because the marginal woman does not use it. Hence the consequences of the new technology upon married women's welfare is intimately related to the situation that prevails on the marriage market.

Finally, men cannot gain from the introduction of the new technology.

When the excess supply of women is large, their utility is not affected. When

the excess supply of women is small, so that the marginal wife does not want a child, the total welfare of the household is increased, but so is the reservation utility of the wife; the husband is left with the same consumption, but loses the benefit he would have received from an unwanted birth. The intermediate case is even more spectacular. Here, all marriages result in a child being born both before and after the reform, so the total surplus generated by marriage is not affected by the innovation. What changes, however, is the intrahousehold allocation of the surplus. The new technology improves the reservation utility of the marginal woman, hence her share of resources; stability requires this shift to take place in all couples. All in all, the new technology results in a net transfer from the husband to the wife, equal to the expected gain of the marginal single woman, i.e. $p(y - z - u(M))$, without any change on the fertility of married couple (who actually do not use the new technology).

We thus conclude that in our model an improvement in the birth control technology, such as the legalization of abortion, generally increases the welfare of *all* women, including those who want a child and are not interested in the new technology. Note, however, that the mechanism generating this gain is largely indirect. The reason why even married women willing to have

a child benefit from the birth control technology is that the latter, by raising the reservation utility of single women, raises the 'price' of all women on the matching market. However, this logic fails to apply in situations of 'large' excess supply of women.

As argued above, empirical evidence suggests that specific submarkets (e.g., low skill African Americans) may exhibit the typical features of a large excess supply of women. In that case, our analysis suggests that married women belonging to this social group may have derived little benefits from the legalization of abortion.

The power shift of the pill The previous argument shows clearly that an important channel through which the new technology benefits all women - what could be called 'female empowerment' - is the raise in their reservation utility. In turn, the source of this change lies in the fact that single women can access the new technology. Consider, now, a situation in which a new technology of this type is introduced but exclusively available to married women. A typical illustration is provided by the pill, which became available first (in 1960) to married women only, and later (by the end of the 1960s) to single women as well (Goldin and Katz 2002).

The case of male excess supply is of little interest, since, in our simple model, all women marry in that case (so that the restriction is irrelevant). On the contrary, in the case of women in excess supply, our previous conclusions are drastically modified by the restrictions on access.

In what follows, let $u(M)$ denote the preference parameter of the marginal woman *in the pre-reform equilibrium*. Also, we first assume that male preference for children is larger than the expected cost of a child - i.e., $u_H \geq p(y - z)$. Then we have the following result:

Proposition 3 *Assume that $u_H \geq p(y - z)$. In the presence of excess supply of women, when the new technology is available to married women only, there exists at least one stable match. For any stable match, all men marry; moreover:*

- *If $u(M) \geq \frac{u}{1-p}$, the new technology is not used and has no impact on the division of surplus in married couples.*

- If $u(M) < \frac{\underline{u}}{1-p}$, there exists two thresholds $u_1(M)$ and $u_2(M)$ satisfying

$$M = \int_0^{u_1(M)} f(t) dt + \int_{u_2(M)}^U f(t) dt,$$

$$u_1(M) \leq u(M) \leq u_2(M)$$

and

$$u_1(M) \leq \underline{u} \leq u_2(M) < \bar{u},$$

which are such that

- Women with a taste parameter $u \geq u_2(M)$ are married and have a child; their welfare is smaller than without the new technology
- Women with a taste parameter $u(M) \leq u < u_2(M)$ are single (whereas they would have been married without the new technology); their welfare is smaller than without the new technology
- Women with a taste parameter $u_1(M) \leq u < u(M)$ are single with and without the new technology, which they cannot access anyway; their welfare is unchanged
- Women with a taste parameter $u < u_1(M)$ are married (whereas they would have been single without the new technology); their

welfare is larger than without the new technology

Finally, men gain from the introduction of the new technology.

In words, making the new technology available to married women only has major, and somewhat counter intuitive consequences. First, the selection of women into marriage is qualitatively different from the initial case, at least when the excess supply of women is not 'large'. Figure 4a gives the new plot of the male maximum utility $\Phi(u)$, in the (interesting) case $u(M) \leq \frac{u}{1-p}$. It thus corresponds, in the new setting, to Figure 3 in the previous context; again, we assume, for expositional convenience, that $p = 0$ for those for whom the pill is available. The key remark is that the graph is no longer monotonic. Indeed, marriage may now also attracts women with a very low taste for children - precisely because they view marriage as the only way to access the birth control technology. Because of this new demand, the competition for marriage becomes tougher; it actually does if $u(M) \leq \frac{u}{1-p}$. Then some intermediate type women are single, whereas they would have been married otherwise. Even for women who are married in both situations, welfare is lower with the new technology, because increased competition on the marriage market reduces their share of total surplus. The gainers from

the new technology are (i) women with the lowest taste for children, and (ii) men.

Insert Figure 4 here

The alternative case $u_H < p(y - z)$ is characterized as follows:

Proposition 4 *Assume that $u_H < p(y - z)$. In the presence of excess supply of women, when the new technology is available to married women only, there exists at least one stable match. For any stable match, all men marry; moreover:*

- *If $M < \int_0^{y-z-\frac{u_H}{p}} f(t) dt$, there exists a threshold $\hat{u}(M)$ such that only women with $u < \hat{u}(M)$ are married. They gain from the technology; no other women gain from the technology; women who were married before the introduction of the technology and are now single lose strictly.*
- *If $\int_0^{y-z-\frac{u_H}{p}} f(t) dt \leq M \leq \int_0^{y-z-\frac{u_H}{p}} f(t) dt + \int_{\hat{u}}^{+\infty} f(t) dt$, then all women with $u < \hat{u}(M)$ are married and gain from the technology; some high type women are married, and lose from the technology; all other women are single and lose from the technology.*

- If $M > \int_0^{y-z-\frac{u_H}{p}} f(t) dt + \int_{\bar{u}}^{+\infty} f(t) dt$, there exists two thresholds $u_1(M)$ and $u_2(M)$ defined by

$$M = \int_0^{u_1(M)} f(t) dt + \int_{u_2(M)}^U f(t) dt$$

and

$$y - z = u_H + pu_1(M) + (1 - p)u_2(M)$$

These thresholds satisfy $u_1(M) \leq u(M) \leq u_2(M)$ and $u_1(M) \leq \underline{u} \leq u_2(M) < \bar{u}$; moreover, they are such that

- Women with a taste parameter $u \geq u_2(M)$ are married and have a child; their welfare is smaller than without the new technology
- Women with a taste parameter $u_1(M) \leq u < u_2(M)$ are single (whereas they would have been married without the new technology); their welfare is smaller than without the new technology
- Women with a taste parameter $u_1(M) \leq u < u_1(M)$ are single with and without the new technology; their welfare is identical
- Women with a taste parameter $u < u_1(M)$ are married (whereas they would have been single without the new technology); their

welfare is larger than without the new technology

Finally, men gain from the introduction of the new technology.

The third case ($M > \int_0^{y-z-\frac{uH}{p}} f(t) dt + \int_{\bar{u}}^{+\infty} f(t) dt$) is illustrated in Figure 4b; it is essentially identical to the previous situation. The first two cases correspond to extreme and somewhat unrealistic situations in which the excess supply of women is so large that only women who do not want children (plus possibly some high type women) get married.

The conclusion of this analysis is clear: reserving the new technology to married women essentially *reverses* the empowerment effect, for two reasons. First, the previous conclusions that all women benefit from the new technology rely on a crucial insight of bargaining model - namely, that the welfare of the marginal ('last single') woman drives the equilibrium allocation. If the new technology is exclusively reserved to married women, this effect disappears, and with it most of the redistribution generated by the innovation. In addition, the availability for marriage of women with a low taste for children, who are willing to accept a lower compensation from the husband for getting married and gaining access to the new technology, toughens competition for

husbands. Women of the high or intermediate type are made worse off by the introduction of the new technology. Only women with a very low taste parameter (i.e., below the lower marginal value) gain from the innovation.

This comparison emphasizes the complex and partly paradoxical welfare impact of a new technology. On the one hand, its effects can go well beyond the individuals who actually use it, or even consider using it. Our model suggests that a major effect of legalizing abortion may have been a shift in the intrahousehold balance of powers and in the resulting allocation of resources, even (and perhaps especially) in couples who were not considering abortion as an option. On the other hand, the new technology benefits all married women only because it is available to singles. A technological improvement which is reserved to married women will have an impact on their fertility, partly because it changes the mechanisms governing selection into marriage. But its impact on women's welfare is either nil or negative, except for a small fraction of women who choose marriage as an access to the new technology.

Our analysis thus suggests that as far as the intrahousehold balance of power is concerned, the introduction of the pill in 1960 (when it was available to married women only) may have had a *disempowerment* effect on most

women. The true empowerment revolution came later; it was caused by the generalization of its availability for single women during the late 60s, and strengthened by the legalization of abortion in the 70s.

Obviously, this analysis is only partial, in that it omits other benefits of the pill (such as women's increased ability to plan their fertility and to achieve higher levels of education). These aspects, which have been intensively discussed in the literature, clearly favored all women, including married one. Still, the strong message of this analysis is that reserving the technology to married women makes an important difference, and, more surprisingly, that most married women are likely to *lose* from this exclusivity.

5 Extensions

5.1 Costly access to the new technology

A first extension of the model introduces the assumption that access to the new technology is 'costly'.²¹ The cost, here, should be understood in a general way; it includes financial costs, but also the moral or ethical discomfort some women may experience with the new technology. The cost, denoted c , is

²¹For expositional clarity, we assume that the innovation drives the probability of unwanted births to zero. Generalization is straightforward.

assumed for the moment to be identical for all women. Since, however, preferences for children differ, a uniform cost generates different responses over the population. Specifically, single women will use the technology only if the net cost of the child (i.e. the income loss minus the benefit u) is large enough to compensate for the abortion cost. Hence:

- a single woman with preference $u \leq \bar{u} - c$ will always abort; the probability of having a child is zero.
- for $\bar{u} - c < u < \bar{u}$, she will neither abort nor try to have a child; the probability of birth is then p
- for $u \geq \bar{u}$, she has a child with probability one.

A similar argument applies to couples, \bar{u} being replaced with \underline{u} .

The resulting equilibrium is depicted in Figure 5. Here, the new male surplus curve (large dots) is located between the pre-reform one (continuous line) and the zero cost one (small dots). The precise computations are straightforward and left to the reader. The general conclusion is that under costly access:

- all women are better off than before the technology became available

(unless the excess supply of women is large, in which case only low type women are better off and the others are indifferent)

- but all women are worse off than if the technology was available at zero cost (except if the excess supply of women is very small, in which case there is no welfare change).

Insert Figure 5 here

A clear policy implication is that *any legislative change that increases the cost of the technology reduces in general welfare of all women, including those who would not use (or even consider using) the technology.* This analysis sheds a new light on the reforms affecting the accessibility to abortion, for instance through Medicaid public funding.²² In the first few years after the national legalization of abortion in 1973, abortion was eligible for Medicaid public funding; this provision was ruled out in 1976 by the Hyde Amendment that generated many similar provisions at the state level. According to our

²²Other examples include the provisions of mandatory counseling and of parental consent for abortion on minors.

analysis, these fluctuations in public funding not only modify women's actual use of abortion, but affect the gains generated by the legalization *for all women* - including those who are willing to bear the costs and those who are not interested in abortion in any case.

5.2 Heterogenous costs

In practice, different women face different costs. The psychological distress clearly varies between women; and even though the financial cost is more uniform, its relative impact on individual consumption and welfare is as heterogenous as individual incomes, a feature that our simplified model does not consider. To capture the idea of heterogenous costs, we make the simple assumption that some women pay the full cost c discussed above, while for others the cost is nil. Also, we assume that the distribution of this cost among women is not correlated with marital status.²³

To further sharpen our point, we consider the extreme case of an infinite cost. Hence, among women with identical preferences some are willing to adopt the new technology while the others do not (or cannot); and the

²³Under the opposite extreme assumption where abortion is costly for single women only, we are back to the case, studied above, of a technology exclusively available to married women.

respective proportion of the two classes (say, μ and $1 - \mu$) is identical for all values of the taste parameter u . Finally, we concentrate on the case in which women are in excess supply. Our results can readily be extended to a less restrictive framework (finite cost, non zero correlation), as well as to the opposite context of male excess supply; this task is left to the reader.

In this new context, women who do not accept the new technology are in a weaker position on the market, because their reservation utility as single is lower. A first consequence is that these women are, everything equal, more likely to marry. Regarding the equilibrium structure, there are now five possible cases. First, when the excess supply of women is 'large' in the previous sense, nothing is changed, since the marginal married woman belongs to the high type and wants children in all situations.

In all other situations, things are more complex. Indeed, married women may either be women with a large taste parameter or women unwilling to adopt the new technology (both being ready to accept a lower compensation). Depending on the size of male supply, four other cases may obtain.

- In the second case, there exist two marginal married women; both belong to the intermediate type, but one accepts the technology while the other does not, the taste parameter being larger for the former. In this

situation, all low type women are single, as well as some intermediate type. This is the case depicted in Figure 6, where the marginal value of the taste parameter before the innovation, $u(M)$, is replaced with the two marginal values $u_N(M)$ (for women refusing the new technology) and $u_A(M)$ (for those accepting it; note that $u_A(M) > u_N(M)$). Stability requires the utility of the respective husbands be the same.

- In the third case, there are still two marginal women, but one (who rejects the technology) belongs to the low type, while the other is 'intermediate and accepts abortion.
- The fourth possible situation is such that all women who reject the technology (including the 'low type' ones) are married, whereas some intermediate women who accept the technology remain single; there is only one marginal woman, who belongs to the intermediate type and accepts abortion.
- Finally, if the supply of male is large enough, the only women remaining single (as well as the marginal married woman) are low type women who accept the technology.

Insert Figure 6 here

From a welfare point of view, it should be noted that no woman can possibly lose from the introduction of the new technology, even when some (possibly many) women reject its use. It is in general the case that all women gain from the introduction, including those who reject it, except when the excess supply of women is large (in which case there is no welfare change). Conversely, if we take as a benchmark the situation in which all women can access the technology, the introduction of costs that preclude its use by some women harms all women in general. Finally, all results are continuous and monotonic in the proportion μ of women willing to adopt the new technology. When μ is close to one, the welfare analysis is largely similar to the initial one; when μ is small, the introduction of the new technology makes little difference altogether.

5.3 Shotgun marriages

In the discussion of the consequences of *Roe vs. Wade*, the issue of 'shotgun marriages' has attracted considerable attention. In an influential paper, Akerlof, Yellen and Katz (1996) have argued that abortion led to the disappear-

ance of shotgun marriages, since women could avoid unwanted pregnancies, a fact that was known to (and potentially used by) men. They conclude that legalizing abortion may actually have harmed some women.

‘There is no such thing as a free marriage’ A first difference between AYK’s model and ours reflects the emphasis we put on the intrahousehold allocation of resources as the endogenous outcome of equilibrium formation on the marriage market. In our model, a woman being ‘shotgun married’ is not necessarily better off than a woman remaining single, especially if the latter can use the new technology. The intuition is simply that in a context of excess supply of women, the surplus generated by marriage will be partly or fully appropriated by the husband; the woman will thus ‘pay’ for marriage by a low share of household consumption.

While this extreme conclusion is clearly linked with our simple setting (e.g., the absence of frictions, hence of post-marital bargaining), we still believe that it stresses an important issue - namely, that shotgun marriage may not come for free. Whether forced marriages closely following an unwanted pregnancy really benefited women is at least debatable; after all, the resulting allocation of household resources was unlikely to favor women, and

it is at least conceivable that the new wife's situation within such 'shotgun couples' was no better than what it would have been had she been single. We believe, in particular, that a crucial and (we think) possibly disputable ingredient of the AYK model is the assumption that the benefits of marriage are exogenously given, and such that women are always better off married than single.

The issue is difficult to assess empirically, if only because modifications of social norms are hard to document (let alone measure). Note, however, that a complete investigation *must* involve estimates of intrahousehold inequality. An analysis that neglects the changes in intrahousehold allocations resulting from the new technology misses a key issue, and may therefore lead to erroneous conclusions. To take but one example, the idea of a 'feminization of poverty' taking place in the 70s should probably be considered with some caution. Insofar as available empirical evidence mostly relies on comparison between individuals and couples while failing to address the crucial issue of the allocation of individual well-being within couples, it may be largely misleading. Standard answers to this problem, based on equivalence scales, are inadequate. Relating the income of a single mother to half (or, for that matter, any fixed fraction) of the couple's income amounts to assuming that

income is split within the couple according to some *exogenously given* rule. Our paper points precisely in the opposite direction: economic theory in general, and equilibrium considerations in particular, tells us that the split is endogenous, driven by the environment, and responsive to the technological changes as well as to market conditions in general. Should these effects be taken into account, the conclusions may be reversed.

In order to establish that some female pauperization resulted from the legalization of abortion, one has to compare the *individual* welfare of women before and after the change. The mere claim, central in AYK's argument, that the reform reduced the number of marriages directly implies that the comparison has to be between the well being of married and single women. From a theoretical perspective, standard, 'unitary' models of household are unable to help, since they do not allow to recover individual utilities within the couple. The alternative, 'collective' approach seems particularly well adapted to tackle the problem. Although available evidence is scarce, Orefice's (2005) findings, already mentioned, suggest that the effects described in our model are well supported by the data. Clearly, additional work on the topic would be welcome. We strongly believe, however, that future progress on this issue will necessarily adopt some variant of the collective model.

Recently, Neal (2004) has provided a detailed discussion of the ‘shotgun’ analysis. Neal raises two questions. If one believes that the decline in shotgun marriages is the primary explanation for the rise in never-married motherhood, why should the impact be concentrated among the economically disadvantaged, and especially among less educated black women? And how come the number of adoptions per non marital birth did not increase over this period? Note that our model provides simple answers for both phenomena. We relate the raise in out-of-wedlock fertility to the sharp decline in men’s supply, a phenomenon specific to the population of black males with lower education. Moreover, in our model the babies born out-of-wedlock were *wanted*, which explains the feeble rate of adoptions.

Did legalization reduce the supply of ‘marriageable’ men? However, a second effect, which is explicitly discussed by AYK and can readily be incorporated into our model, may reverse our conclusion. Specifically, assume, following AYK, (i) that a fraction of the male population is not interested in marriage, but only in sexual activity, (ii) that in the absence of the birth control technology, social norms impose an implicit commitment from the male part, whereby sexual activity leading to pregnancy must end up in marriage,

even against the male's initial intention, and (iii) that the availability of abortion results in the disappearance of the social norm, in the sense that the father of an unwanted child no longer feels committed by the mother's decision to keep the child. Under such circumstances, the new technology may result in a reduction in male supply. According to the previous analysis, such a reduction may be harmful for women in two cases: when the initial equilibrium was characterized by an intermediate excess supply of women, and when the type of equilibrium changes because of the cut in male supply.

This interpretation has empirical implications that can be tested. In our framework, the 'impoverishment of women' effect described by AYK must operate through a particular channel; namely, a significant fraction of the male population decides to remain single, while they would have chosen (or be forced) to get married before legalization. One way to empirically analyze this issue, hence, is to see whether legalization had a significant impact on the probability that a male will remain single.

Casual observation does not seem to support this prediction. The proportion of single men in the male population does not seem to respond to legalization (see Figure 7).

Insert Figure 7 here

To formally confirm these observations, we regress the probability of being single on age, education, race and fixed effect by year and state, as well as an abortion dummy on the male population aged 15-50 in the CPS March supplements 1968-1980. We find that the abortion dummy is not significant, and its point estimate is actually negative. The results seem robust to various specifications.

These findings suggest that the central mechanism in AYK, namely that a significant proportion of males decided to withdraw from the marriage market (hence remain single) following legalization, does not seem to be supported by the data. Clearly, though, this conclusion is preliminary; additional empirical work on this topic will be needed.

6 Final comments

The model proposed in this paper provides a simplified view of the phenomena at stake. Many issues remain open. Marriage markets are characterized by multidimensional heterogeneity (tastes, but also incomes, ...).

Frictions are paramount, which implies that intrahousehold bargaining could profitably be taken into account.²⁴ Dynamic issues are crucial, not only because divorce and remarriage are important features of the market, but also because such factors as the average age at marriage or the age difference between the spouses are known to matter for equilibrium. A single, unified market for marriage probably does not exist; the relevant concept is more a multiplicity of markets by age, localization, race and religion. Our model applies to each of these submarkets, with possibly different conclusions in the various situations. For instance, the decrease in the supply of men may be a more stringent phenomenon in some contexts (say, younger population in impoverished urban neighborhoods) than in others - and this fact may explain considerable differences in marriage or out-of-wedlock fertility across the population. Clearly, the empirical definition of these submarkets and the measure of their interconnections will raise delicate problems for future research.

Homogeneous males We may also briefly come back to our basic assumptions. First, the assumption that men are identical is essentially harmless.

²⁴This view is explored by Chiappori and Weiss (2003, 2005) and Chiappori, Iyigun and Weiss (2005).

Heterogeneity in male income would not change our qualitative conclusions. When women are relatively more numerous, our quasi linear setting implies that male heterogeneity is irrelevant from women's perspective, because it affects neither the *surplus* generated by marriage nor its distribution. When males are in excess supply, while income matters (e.g., wealthier males are more likely to get married), the qualitative conclusions of the model remain unchanged. The case of heterogeneous male preferences for children is slightly more difficult. One can readily check that in our context, the surplus generated by marriage is supermodular, leading to assortative matching. The distribution of this surplus between spouses is more complex, and depends on the respective densities of the male and female distributions of preferences. However, our qualitative conclusions remain valid. When women are in excess supply, the welfare of each wife still crucially depends on the situation of the marginal woman (i.e. of the 'last' single), which leads to similar predictions.

Commitment versus bargaining A second, crucial ingredient of our model is the commitment assumption. Our conclusions clearly rely on the fact that initial conditions (i.e., the state of the 'marriage market' and the

legal environment) have lasting effects on the allocation of resources within marriage. The commitment assumption is, if anything, dominant in the literature. For instance, all models of the marriage market based on matching or search theory rely on the same commitment assumptions as our.

Still, whether commitment is feasible in practice - i.e., whether initial agreements can be enforced in the medium or long run - is a legitimate question. One way to implement commitment is through prenuptial contracts; after all, legally enforceable contracts are a standard way of achieving commitment in social life, and marriage should be no exception. We believe, however, that commitment is, if anything, easier to achieve in a marriage context, because of the non-monetary aspects involved. Spouses typically have strong attachment to each other as individuals; this affection can be expected to increase the psychological cost of renegeing, thus stabilizing the initial agreement. We would actually submit that the empirical scarcity of prenuptial agreements, despite their general availability and their relatively low cost, suggests, if anything, that enforcing commitment without contracts may not be such a difficult task in most situations.

It is interesting to compare our conclusion with those stemming from a model based on spot bargaining, whereby intrahousehold allocations are

constantly renegotiated. In both contexts, abortion would favor women who are willing to use it. However, in a bargaining context, the existence of abortion as a possible outcome ameliorates women's bargaining position *only insofar as the threat that they may use it is credible*. Most of the time, this is hardly the case: if most women care more about children than men do (a standard assumption in the literature), abortion is unlikely to tilt significantly the intrahousehold balance of power. Moreover, for those couples in which abortion is viewed by the wife as an option, bargaining positions are affected irrespective of the situation of single women, so restricting access to married women should make little difference. In our model, the opposite logic prevails: the situation of single women is crucial, and it affects the situation of all women, including those who would not even consider abortion as a possibility.

In the end, the empirical relevance of the commitment assumption should be viewed as an empirical issue. From this perspective, a contribution of our model is precisely to point out the differences between the commitment and bargaining approaches and the contrasted testable predictions they generate.

In conclusion, our approach puts forth important insights regarding the

impact of birth control technologies. Our main message is that issues related to intrahousehold allocation are of crucial importance, and that the distribution of resources and welfare within couples must be understood as an *endogenous* phenomenon which responds to changes in the economic and technological environment. An important consequence is that the introduction of new technologies (or new benefits, for that matter) may have a major impact on individual welfare *even within couples who do not directly benefit from it*. This intuition has been frequently mentioned by sociologists (the idea of ‘female empowerment’ resulting from the legalization is an old theme of feminist studies), but it may have been somehow disregarded by economists, at least as far as explicit modeling is concerned.²⁵

Economic analysis provides new and interesting insights on the ‘female empowerment’ issue. For instance, the impact of a given reform deeply varies with such exterior determinants as female wages, the generosity of benefit systems or the situation on the market for marriage. These interactions are complex, and we believe that any empirical assessment would greatly benefit from a precise theoretical investigation of the main mechanisms at stake. To give an example, our model suggests that there is no simple relationship

²⁵ An interesting and early exception is the analysis of guaranteed employment programs in India proposed by Haddad and Kanbur (1990).

between the extent an innovation is used and the magnitude of its impact on intrahousehold inequality.²⁶ Studies analyzing the fertility consequences of the innovation, although important, are no substitutes for an investigation of intrahousehold allocation issues. This is a task for which specific models must be devised.

A final contribution of our model is precisely to show how the ‘collective’ analysis, which has been largely developed at the household level, can easily be incorporated into a market-wide analysis to provide a model of the type required. In a sense, the link is very natural. The ‘sharing rule’, a crucial ingredient of the collective approach, has a direct translation in the context of a matching model; specifically, it defines the imputations that stabilize a given match. Because our approach is based on matching, it seems well adapted to analyze the impact of the reform on couples formed *after* the reform has been implemented; in that sense, it complements alternative works based on bargaining models.²⁷ Our contribution, in this sense, should be seen

²⁶Assume a given innovation (say, abortion legalization) is simultaneously introduced in two different ‘submarkets’, respectively characterized by a ‘small’ and a ‘large’ excess supply of women (in the sense defined in the paper). The consequences on fertility will be much larger in the second submarket (all single women of ‘low’ and ‘intermediate’ type will use the technology), whereas the consequences on intrahousehold allocation is nil. Conversely, in the submarket with small excess supply, few women will use the technology (only those belonging to the low type), but the allocative impact will be maximum.

²⁷See for instance Chiappori and Donni (2005)

in the perspective of a general line of research which is currently pursued.

References

- [1] Aiyagari, S. Rao , Jeremy Greenwood and Nezih Guner, ‘On the State of the Union’, *Journal of Political Economy*, 108, 2 (2000), 213-244
- [2] Akerlof, G., Yellen J., Katz M. (1996), “An analysis of out-of-wedlock childbearing in the United States”, *Quarterly Journal of Economics*, 111- 2, 277-317
- [3] Ananat, Oltmans E., J. Gruber and P. B. Levine (2004), ‘Abortion Legalization and Lifecycle Fertility’, *Working Paper 10705, NBER*
- [4] Angrist, J. (2002), ”How do sex ratios affect marriage and labor markets? Evidence from America’s second generation”, *Quarterly Journal of Economics*, 117-3, 997-1038
- [5] Angrist, J., Evans, W. (1996), “Schooling and labor market consequences of the 1970 state abortion reforms”, *Research in Labor Economics*, 18
- [6] Gary S. Becker, ‘A Theory of Marriage: Part I’, *Journal of Political Economy*, 81-4, 813-846

- [7] Gary S. Becker, 'A Theory of Marriage: Part II', *Journal of Political Economy*, 82-2, S11-26
- [8] Gary S. Becker, *A Treatise on the Family*, Harvard University Press, 1993
- [9] Browning, M. and P.-A. Chiappori (1998), "Efficient Intra-Household Allocations: a General Characterization and Empirical Tests", *Econometrica*, 66, 1241-1278.
- [10] Browning M., Chiappori, P.-A., Weiss, Y. (2005), "A simple matching model of the marriage market", *mimeo*, University of Chicago
- [11] Burdett, Kenneth, and Melvyn G. Coles, 'Marriage and Class'. *The Quarterly Journal of Economics*, 112, 1 (1997), 141-168
- [12] Burdett, Kenneth, and Melvyn G. Coles, 'Long-Term Partnership Formation: Marriage and Employment', *The Economic Journal* 109, 456 (1999), F307-F334
- [13] Chiappori, P.A, M. Iyigun and Y. Weiss (2006), "Spousal Matching, Marriage Contracts and Property Division in Divorce", *mimeo*, Columbia University.

- [14] Chiappori, P.-A., Weiss, Y. (2002), "Marriage contracts and divorce: an equilibrium analysis", *mimeo*, University of Chicago
- [15] Chiappori, P.-A., Weiss, Y. (2004), "Divorce, remarriage and child support", *mimeo*, Tel Aviv University
- [16] Choo E. and Siow, A. (2003), "Who marries whom and why", *Journal of Political Economy*, 114, 1 (2006), 175-201
- [17] Coombs, L., Fernandez, D. (1978), "Husband-wife agreement about reproductive goals", *Demography*, 15-1, 57-73
- [18] Edlund L. (2004), "Sex and the city", *Scandinavian Journal of Economics*, 107-1, 25-44
- [19] Goldin, C., Katz, L. (2000), "Career and marriage in the age of the pill", *American Economic Review*, 90-2, 461-465
- [20] Goldin, C., Katz, L. (2002), "The power of the pill: oral contraceptives and women's career and marriage decisions", *Journal of Political Economy*, 110-4, 730-770
- [21] Haddad, L., Kanbur, R. (1990), "How serious is the neglect of intra-household inequality?", *Economic Journal*, 100-402, 866-881

- [22] Haddad, L., Kanbur, R. (1992), "Intra-household inequality and the theory of targetting", *European Economic Review*, 36, 372 - 378
- [23] H eritier, F. (2002), "Masculin/F eminin II", *Odile Jacob*
- [24] Iyigun, M., Walsh, R. (2004), "Building the family nest: a collective household model with competing pre-marital investments and spousal matching", *mimeo*, University of Colorado
- [25] Klerman, J. (1999), "US Abortion Policy and Fertility", *American Economic Review*, 89-2, 261-64
- [26] Levine, P. (2004), "Sex and consequences", *Princeton University Press*
- [27] Levine, P., Staiger, D., Kane T., Zimmerman, D. (1999), "Roe vs. Wade and American fertility", *American Journal of Public Health*, 89-2, 199-203
- [28] Lundberg, Shelly , and Robert A. Pollak, 'Bargaining and Distribution in Marriage', *The Journal of Economic Perspectives*, 10, 4 (1996), 139-158

- [29] Michael, R., (2000), "Abortion decisions in the U.S", in Laumann, E, Michael, R (eds.), *Sex, Love and Health: Public and Private Policy*, University of Chicago Press
- [30] Moffitt, R. (1992), "Incentive effects of the US welfare system: a review", *Journal of Economic Literature*, 30-1, 1-61
- [31] Mortensen, Dale (1988), "Matching: Finding a Partner for Life or Otherwise," *American Journal of Sociology*, 94 (supplement), s215-s240.
- [32] Neal, D. (2004), "The relationship between marriage market prospects and never-married motherhood", *Journal of Human Resources*, 39-4, 938-957
- [33] Oreffice, S. (2005), 'Did the legalization of abortion increase women's household bargaining power? Evidence from labor supply', mimeo, University of Chicago.
- [34] Alvin E. Roth and Marilda Oliveira Sotomayor (1990), *Two-Sided Matching*, Econometric Society Monograph, Cambridge University Press.

- [35] Shimer, Robert , and Lones Smith, ‘Assortative Matching and Search’, *Econometrica*, 68, 2 (2000), 343-369
- [36] Siow A. (2002), “Do innovations in birth control technology increase the welfare of women?”, *mimeo*, University of Toronto
- [37] Western, B., Pettit, B. (2000), "Incarceration and racial inequality in men’s employment", *Industrial and Labor Relations Review*, 54, 3-16
- [38] Williams, L. (1994), “Determinants of couple agreement in US fertility decisions”, *Family planning perspectives*, 26-4, 169-173
- [39] Wilson, W. (1987), ”The truly disadvantaged: the inner city, the underclass, and public policy”, *University of Chicago Press*

Table 1
Agents' utilities in the various equilibria.

	ESW			ESM
	Large	Intermediate	Small	
Marriage	$u(M) \geq \bar{u} = y - z$	$u \geq u(M) \in (\underline{u}, \bar{u})$	$u(M) \leq \underline{u} = y - z - u_H$	All women married
H women	$z + u$	$(y - u(M))(1 - p) + pz + u$	$z + (1 - p)u_H + u$	$z + u_H + u$
I women	$(1 - p)y + p(z + u)$	married: $(y - u(M))(1 - p) + pz + u$ single: $(1 - p)y + p(z + u)$	$z + (1 - p)u_H + u$	$z + u_H + u$
L women	$(1 - p)y + p(z + u)$	$(1 - p)y + p(z + u)$	$(1 - p)y + p(z + u)$	$(1 - p)y + p(z + u + u_H)$
Men	$Y + u_H$	$Y - (y - z - u(M))(1 - p) + u_H$	$Y + pu_H$	Y

Notations: 'Marriage' denotes the threshold on the parameter taste u above which marriage takes place.

ESW: excess supply of women; ESM: excess supply of men.




Figure 1
Maximum husband's utility as a function of the wife's taste
Initial model




Figure 2a
Increasing female income
 (post-increase in dotted line)




Figure 2b
Reducing male supply (intermediate ESW)
 (post-increase in dotted line)

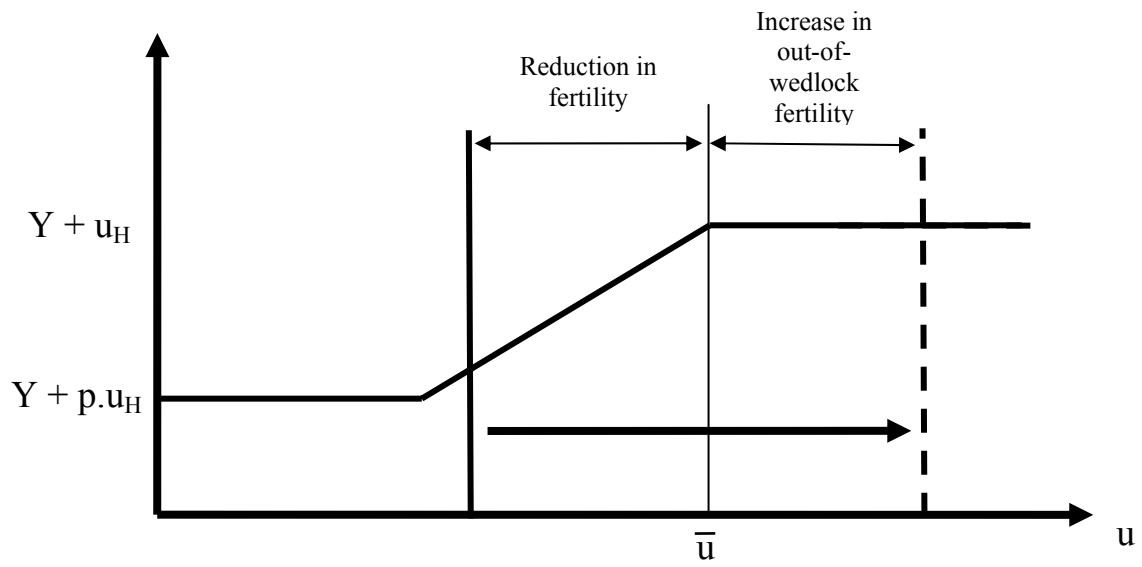


Figure 2c
Reducing male supply (from intermediate to large ESW)
 (post-increase in dotted line)




Figure 3
The impact of the new technology





Figure 4a
Maximum husband's utility as a function of the wife's taste
Innovation available to married women only; case $u_H > p(y-z)$

Figure 4b

**Maximum husband's utility as a function of the wife's taste
Innovation available to married women only; case $u_H < p(y-z)$**






Figure 5
Costly access to the new technology




Figure 6
Heterogeneous costs
 (women who accept the technology are represented by the dotted line)




Figure 7a

Percentage of singles in the male population, age 14 and above

Source: Tables 1 of the annual March Census Current Population Reports "Marital Status and living arrangements", issues from March 1968 to March 1979"




Figure 7b

Percentage of singles in the male population, age 25-29

Source: Tables 1 of the annual March Census Current Population Reports "Marital Status and living arrangements", issues from March 1968 to March 1979"