Family Size and the Demand for Sex Selection: Evidence From China

Samuel Marden*
University of Sussex

March 10, 2016

Abstract

In China, many fewer girls are born than would be expected given natural birth rates. This imbalance has worsened dramatically over the last 40 years. The roughly contemporaneous fall in fertility per woman is often mooted as a source of this apparent increased demand for sex selection: fewer births make it harder to have a son by chance. Despite this, causal evidence is limited. This paper exploits geographic variation in changes in fertility, arising as a consequence of China's agricultural reforms (1978-84), to provide this evidence. Specifically, I show that households living in counties that benefitted more from the reforms, increased their fertility relative to households elsewhere. I then show that these households are also less likely to engage in sex selection. These changes appear to have been due to higher local incomes interacting with the enforcement of the One Child Policy. The timing of the changes in fertility and sex selection are informative: while fertility increased almost immediately, the decline in sex selection only emerged from the mid 1980s—contemporaneous with the widespread availability of ultrasound. These results suggest that the dramatic decline in fertility in 1970s China, as well as the smaller decline due to the One Child Policy in the 1980s, may have had an important role in fuelling the demand for sex selection.
JEL Codes: J11, J13, J16

*s.marden@sussex.ac.uk. I gratefully acknowledge financial support from the Economic and Social Research Council and the Royal Economic Society.

1 Introduction

Many countries, particularly in South and East Asia, have shares of men in the population that far exceed those observed in Western Europe and North America. This surfeit of men implies a shortfall of women, the *missing women* of Sen (1990, 1992). Although these women are mostly missing from lower and middle income countries, development does not appear to resolve the imbalance. Rapid economic growth over the past 30 years in India, China, South Korea and Taiwan has been accompanied by increasingly biased sex ratios.

Several factors have been implicated in this trend. Parents in South and East Asian countries often have a preference for sons, which can lead to sex selective abortions and differential levels of care during childhood.[1] It is thus possible that the rapid economic and social changes in these countries may have strengthened preferences for sons.[2] Regardless of whether preferences changed, widespread availability of ultrasound technology from the early 1980s undoubtedly reduced the cost of acting on these preferences and led to a boom in sex selective abortions (Chen et al., 2013; Lin et al., 2014). Finally, by reducing the probability of having a boy by chance, fertility transitions in these countries may have increased the number of parents having one by design (Park and Cho, 1995; Das Gupta and Mari Bhat, 1997; Das Gupta and Shuzhuo, 1999).

Although the role of ultrasound technology seems clear, untangling the impact of fertility transition from preference change and economic transformation has proven challenging. Jayachandran (2014) provides compelling evidence that lower hypothetical fertility increases the *desired* share of boys, but it is not obvious how desire translates into action. Perhaps then, the most persuasive evidence in favour of the fertility transition hypothesis comes from two papers evaluating the consequences of the One Child Policy (Ebenstein, 2010; Bulte et al., 2011). These papers, which respectively exploit provincial variation in financial penalties and the differential enforcement faced by minorities, show that more restrictive policy is associated with lower fertility and more biased sex ratios. However, because of the strong regional component in the variation exploited by Ebenstein and the cultural differences between the treatment and control group of Bulte et al., it is not clear whether these changes are driven by pre-existing differences in preferences interacting with the arrival of sex selective abortion, differential trends in preferences for total fertility and sons, or the fertility restrictions imposed by the One Child Policy (1CP).

The key empirical challenge in understanding the interaction between total fertility and sex selection, is that changes in desired and realised fertility are generally endogenous to preferences for sex selection. This paper overcomes this endogeneity, by exploiting variation in fertility induced by differential exposure to the benefits of China's agricultural reforms (1978-84). In particular, it exploits the fact that parts of China suited to cash crops—here cotton and oilseeds, the two most important cash crops at the time—had significantly faster post reforms growth in both agricultural and non-agricultural output (Lardy, 1983; Marden, 2015). Expanded freedom to grow cotton and oilseeds increased agricultural output by around 10% nationwide between

[1] In fact, as shown by Anderson and Ray (2010), women face excess mortality throughout the life cycle.

[2] One channel that economic outcomes have been shown to affect son preference is through changes in women's labour market prospects (Rosenzweig and Schultz, 1982; Qian, 2008).

1978-85, but much more in counties endowed with land suited to their production. Because the economic gains accumulated rapidly, and underlying preferences change only slowly, we can be confident that changes in sex selection are not the result of preference change.[3]

I show that places that benefitted more economically from the reforms had higher subsequent fertility. Using a difference-in-difference strategy, and the fertility histories of more than 700,000 married women born between 1945 and 1960 (13.7 million woman-year observations), I show that women resident in counties of China endowed with land suited for growing cash crops—which had higher post reform incomes—had significantly higher fertility after the reforms than their counterparts residing in counties suited to grain. These results exploit only within province-variation and hence compare households that are likely to have similar underlying preferences for total fertility and gender. I also include a full set of woman specific fixed effects to sweep out time invariant differences between households. Consistent with the key parallel trends assumption, there were no differential trends in fertility prior to the reforms. The results indicate that the income elasticity of demand for children was positive in the wake of the reforms.

Next, I show that that households which benefited more from the reforms and increased their fertility also engaged in less sex selection.[4] Exploiting the fact that, for second born children, households only appear to engage in sex selection if the firstborn child is a girl (Figure 1; Yi et al., 1993; Ebenstein, 2007), I use a triple-difference strategy to investigate the effect of higher fertility on sex ratios. The differences are: (1) the-timing of the reforms; (2) how beneficial the reforms were; and (3) the gender of the older sibling. This formulation allows for pre-existing differences in son preferences by county, county specific economic and environmental shocks to 'natural' gender specific birth and survival rates, and differential trends in sex selection over time (by province). As with fertility, there was no differential trend in sex selection prior to the reforms, whereas after the reforms, the relative level of sex selection declined in the areas that had relative increases in fertility.

Interestingly, while observed fertility increases even as the reforms are ongoing, from the early 1980s, the decline in sex selection apparent from the mid 1980s onwards. The timing of the relative decline in sex selection is contemporaneous with the widespread availability of cheap sex detection technology. Higher fertility may have reduced the underlying demand for sex selection, but this lower demand only became relevant when the cost was low enough for sex selection to be a feasible option. This suggests that the dramatic falls in fertility during the 1970s may have been an important factor underpinning the demand for sex selection apparent in the 1980s.

The results in this paper suggest that the economic shock I exploit was associated with a positive income elasticity of demand for fertility, and a strongly positive income elasticity of demand for girls. However, under normal circumstances higher incomes

[3]For evidence on the slow pace of changes in son preference, see Almond and Edlund (2008) and Abrevaya et al. (2009) who provide evidence of sex selection by Asian households resident in the US, and Dubuc and Coleman (2007) who provides similar evidence for the UK.

[4]In fact, what I show is that the share of women in a given cohort that have survived until 1990 is higher in places relatively suited to cash crops after the reform than before the reform. However, as there is limited gender biased excess mortality for children in China (Anderson and Ray, 2010), these differences primarily reflect prenatal sex selection and excess mortality immediately following birth.

are associated with lower fertility. Understanding why, in this case, a positive effect on fertility is observed is important for the interpretation of the results on sex selection.

One possible explanation is that higher incomes interact with the enforcement of the One Child Policy (1CP). Since 1979, the 1CP has limited fertility by imposing sanctions on households which breach prescribed family sizes.[5] Because the sanctions are primarily economic, the relatively well off have often been willing and able to bear these sanctions (Scharping, 2013, pp. 142-143). From the early 1980s, the severity of the fines was set at the provincial level, so within a province, richer individuals may be less constrained in their fertility by the 1CP.[6] If this were the case then, other things equal, richer households would have more children in general and, under the fertility transition hypothesis, more girls in particular.

To test this, I exploit the fact that some households are less constrained by the 1CP than others. Specifically, under the 1CP, minority households were typically permitted to have more than the 1.5 children than Han households were nominally limited to. Comparing the effects on fertility and sex selection on Han and non-Han households reveals that the observed effects are solely due to the effect on Han households—the households most affected by the 1CP. This is consistent with higher incomes reducing the bindingness of the 1CP, either by making the fines more affordable, by inducing laxer local enforcement, or by otherwise increasing the willingness to pay for an additional child.[7] These results are consistent with those of Anukriti (2014), which shows that financial incentives for lower fertility can lead to more biased sex ratios, even when additional incentives for girls are provided.

An alternative explanation for the change in sex ratios, if not for the increase in fertility, is that of Qian (2008), which, using an empirical strategy similar to that employed in this paper, showed that changes in the (relative) economic status of women in China could affect women's outcomes. The fact that the effect observed in this paper appears to be a consequence of higher local incomes interacting with the 1CP, is suggestive that changes in relative labour market outcomes are not driving the results. Furthermore, although women form a somewhat larger share of the labour force for cotton production than farming in the population (data are not available for oilseeds), the ability to grow cotton and oilseeds on a large scale has no observable effect on women's labour market outcomes. Women do not appear to have gained disproportionate economic benefits from growing these cash crops.

This paper is most closely related to Almond et al. (2013), which also explores the

[5]For most households, most of the time, the 1CP does not actually prescribe a limit of just one child under all circumstances. For instance, most rural Han Chinese households were permitted an additional child if the first child was a daughter. Minority households were often allowed two or sometime even three children.

[6]This is not to imply that fines were invariant within a province. Indeed, they were typically related to both household income (albeit imperfectly so in the countryside where measurement was highly problematic) and place of residence. Furthermore, the enthusiasm with which they were enforced also varied. Nevertheless, with decreasing marginal utility of income, even a perfectly proportional fine is less onerous for a rich household than a poor one.

[7]As labour utilisation was typically higher for cash crops than for grains (Taylor, 1988) is also possible that growing cash crops increased the demand for children as a source of labour. However, because of the relatively high ratios of land to labour, and the large quantities of 'surplus' adult labour created by the reforms it is unlikely that households faced a major labour shortage. Consistent with this, Bowlus and Sicular (2003) document that, in the early 1990s, an additional child increases total household agricultural labour supply by just 25 days.

affect of agricultural reform related increases in income on sex selection in second born children. However, despite the similarities the results are quite different. While I find higher incomes increased fertility and reduced sex selection (consistent with higher incomes alleviating the fertility constraints imposed by the 1CP), Almond et al. find that China's higher incomes *increased* sex selection. This difference is likely attributable to two factors. First, rather than exploit cross sectional variation in the suitability of the reforms interacted with the general timing of the reforms, they exploit variation in the specific timing of the rollout of the reforms at the county level. As the rollout was completed almost entirely over the years 1982-1984, their strategy is designed to capture vary short run effects of higher incomes. Second, unlike the present paper, they do not allow for differential trends in sex selection conditional on the gender of an older sibling. As sex selection increased substantially after the reforms for households with daughters but no sons, it is possible that this omission biased their results towards finding an increase in sex selection.[8]

This paper is also closely related to Qian (2008), which uses the increased economic returns associated with tea and orchard crops in the wake of China's reforms to explore the consequences of the differential labour market performance of men and women on gender bias. She shows that female (male) biased income changes reduce (increase) son preference. This paper asks a different question: it uses variation incomes provided by the reforms to show that the economic sanctions associated with the One Child Policy did constrain fertility. Then using this plausibly exogenous variation in fertility, it explores the link between fertility transition and sex selection which has often been argued to have been a major force in much of Asia's increasingly imbalanced sex ratio.

The remainder of the paper proceeds as follows. Section 2 provides additional background to sex selection, fertility policy and agricultural reform in China. Section 3 provides a simple conceptual framework for thinking about the interaction between income, fertility and sex selection in the presence of son preference. Section 4 provides the empirics, first establishing the link between suitability and incomes (Section 4.1), then exploring the link between incomes and fertility (Section 4.2), before finally evaluating the link between fertility and gender ratios (Section 4.3). Section 5 concludes.

2 Background

2.1 Fertility and Fertility Policy in the Peoples Republic of China (1949—)

This section provides an introduction to fertility and fertility policy in the Peoples Republic of China (PRC). It is based on the much more extensive treatments of the topic provided by Chang et al. (2005), White (2006) and Scharping (2013).

[8] Their most rigourous specification is $Boy^2_{ijt} = \alpha + \beta_2 Girl^1_{int} + \beta_3 Reform \times Girl^1_{ijt} + \gamma_{jt} + \epsilon_{ijt}$ where the unit of observation is the second born child, *Boy* is a dummy for if it is a boy, *Girl* a dummy for if it's older sibling was a girl and γ_{jt} and county-by-time fixed effect. (Almond et al., 2013, Equation 2, p.11). The citation provided is for the 2013 NBER working paper version, which is likely to be stable and consistently available, however a version dated May 2014 employs the same specification.

2.1.1 Early Years to 1978

In the early years of the PRC little was done to control the population. Tentative steps were made to increase the availability of contraceptives, first in urban areas and then in the countryside. These initial efforts, expanded into campaigns encouraging later marriage and longer gaps between births. However, these early forays into population control were soon abandoned in the wake of the collapse in birth rates and chaos accompanying the disastrous Great Leap Forward (1959-61).

Fertility rebounded after the end of the Great Leap Forward. Birth control programs were quickly reinstated in cities, but little action was taken in the countryside. Total fertility fluctuated at around 6 births per woman for the rest of the sixties.[9] Concerns over high birth rates eventually led to fertility control policies being reintroduced to the countryside. Campaigns exhorting that 'one [child] is not too few, two, just right, three too many' and that births ought to be 'later, longer, and fewer', where longer refers to gaps between births, were instrumental in the precipitous decline in fertility from more than 6 births per woman in the late 1960s to 2.7 in 1978.

2.1.2 From 1979: The One Child Policy

The 'One Child Policy' (1CP) was introduced beginning in 1979. In the early days, the degree to which it was enforced varied widely. Sanctions were sometimes draconian and in some areas included forced abortions, sterilisations and severe financial penalties. Passive resistance to the 1CP in the countryside was widespread.

Faced with this resistance, the one child policy was somewhat standardised beginning in 1984. From this time, most rural households were allowed an additional child if their first child was a daughter—the 1.5 child policy—and a number of other exemptions were carved out. Minorities were also brought into the 1CP, although they were usually allowed a second, and sometimes even third, child. Despite this liberalisation, the 1CP appears to have been successful in further reducing total fertility in China—to 2.3 births per woman in 1990 and 1.5 in 1998 .

Throughout its existence, the 1CP has been enforced (although not exclusively so) by the imposition of fines and 'fees'.[10] After the initial period, guidance for the level of fines has been set at the provincial level, and fines have nominally been proportional to household income (although due to the challenges associated with accurately measuring income, fines were increasingly standardised). Ebenstein (2010) has shown that, at the province level, the increase in sex selection was inversely proportional to the severity of fines, suggesting variation in the burden of the fines had real consequences for fertility. Anecdotally, the relatively well off appear to have been willing to pay the fines (Scharping, 2013, pp. 142-143) which suggests that holding the regime of fines and sanctions roughly constant—for instance, within a province—incomes may interact with the effectiveness of the 1CP. The evidence in this paper will indicate that, within a province, higher incomes were associated with higher fertility, which is consistent with this anecdotal evidence. However, given the latitude that local authorities had in the

[9] Total fertility is the expected lifetime fertility per woman.

[10] Other sanctions have included, non-adjustment of land/living space for increased family size, exclusion from health, welfare and other social services, reduced rations, and expulsion from the communist party.

severity of enforcement, even after the standardisation of the 1CP, it is also possible that economic success increased local government scope for laxer enforcement.

2.2 Sex Ratios in China

While Chinese fertility remained high, sex ratios for newborns and young children remained close to the natural rate of around 1.06 boys for every girl. In the 1953 and 1964 censuses, there were 1.05 and 1.04 boys age below 1 for every girl. By the early 1980s there were 1.08-1.09 boys, by the end of the 1980s, 1.11, and by 1995, 1.16 (figures from Scharping, 2013, Table 32). Given that fertility declined dramatically in the 1970s but sex ratios at birth barely moved, it is hard to make the case that fertility transition described above was the sole cause of these changes. In fact the timing of the increase suggests that the availability of sex detection technology may have been the proximate cause for the increase (Chen et al., 2013; Lin et al., 2014). This paper will ask whether the change in fertility during the 1970s, and to a lesser extent through the 1980s and 1990s were instrumental in generating the required latent demand for sex selection.

As documented by Yi et al. (1993), Ebenstein (2007) and others, these biased sex ratios are driven *entirely* by gender imbalance in higher order births. Figures 1-3 provide a demonstration of this. Using unadjusted data from the 1990 census (so these data also include differential mortality up to age '1990 minus birth year'), and conditioning on the gender of older siblings, I have plotted the share of girls in the population, in a given birth cohort, conditional on being the firstborn child (figure 1a); second born child by gender of older sibling (figure 1b); and third born child by gender of older siblings (figure 1c). For first births, second births to families with a boy, and third births to families with a boy and a girl, the share of girls is extremely close to the natural share. However, for families with a single girl, or two girls, a substantial and growing bias in favour of boys is immediately apparent. Interestingly households with two boys have a significant bias in favour of girls. However, because more families choose to have second and third children if they have no boys than if they have no girls, the overall effect of the sex selection is to strongly increase the share of boys in the population.

2.3 Agricultural Reforms, Incomes and the 1CP

This paper will exploit a plausibly exogenous shock to rural incomes to provide variation in fertility. Because the shock to incomes is plausibly unrelated to underlying preferences for fertility and sons, this variation avoids the usual endogeneity preferences to fertility rates that have limited the scope for identifying a causal link between fertility and sex selection. Furthermore, because the increase in fertility will appear to be the result of an interaction between higher incomes and the 1CP, and that the income shock does not affect gender specific labour market outcomes, we can be confident that the link is not due to the increased economic returns to having girls in these areas. Because child labour utilisation in China was low (Bowlus and Sicular, 2003), and rural China had large quantities of 'surplus' labour at this time, we can also be confident that the increase in fertility was not due to an increase in the demand for labour in the household.

Variation in incomes is obtained by exploiting features of China's agricultural reforms (late 1978-84). These reforms marked the start of the reform-era and resulted in

substantial increases in agricultural productivity and output. Their success paved the way for non-agricultural reform, and the economic successes of the next three decades. The reforms sharpened incentives, by decommunalising agriculture and increasing the prices faced by farmers (McMillan et al., 1989; Lin, 1992), and improved the allocation of crops to land, by *de facto* liberalising planting decisions (Lardy, 1983; Marden, 2015).[11]

This paper will exploit the fact that the gains from the *de facto* liberalisation of planting decisions were not distributed uniformly across china. Pre-reform institutions heavily encouraged the production of grain over other crops in three main ways. First, rural areas were only allowed 'to produce economic crops or raise animals, [...] only after they had achieved basic self-sufficiency in food grains' (Lardy, 1983, p. 49). Self sufficiency was enforced through the state's monopoly on trade in agricultural produce and compounded by the low productivity of Chinese agriculture which kept most rural households close to subsistence. Second, Communes and production teams were also often required to deliver quotas of grain to the state. Third, the political success of rural party cadres was linked to the production of grain. The reforms relaxed these constraints, and allowed farmers who wished to grow cash crops to somewhat specialise in their production.

Thus, China's agricultural reforms reduced the emphasis on grain production and sharpened the incentives faced by farmers. Agricultural output duly expanded rapidly. Grain output increased by 5 percent per-year between 1978 and 1985, compared to 2.4 percent per-year between 1952 and 1978. Output of cash crops increased even more rapidly: cotton by 19.2 percent per-year, sugar by 12.3 percent and oil-crops by 14.8 percent, compared to 2, 4.5 and 0.8 percent per year pre-reform. Because farmers with land suited to cash crops were best placed to benefit from planting it, agricultural output grew fastest in areas suited to cotton and oilseeds (Marden, 2015).[12]

3 Conceptual Framework

To motivate the empirical analysis, I provide a simple economic model of fertility at the household level in the tradition of Becker (1960). Households consist of a man and a woman who maximise their joint utility (other household types are ignored for simplicity). Households have preferences over consumption $u(c)$, where u is strictly increasing and strictly concave, and family size n. In addition to these standard elements, households also prefer to have at least one son and receive a payoff of γ_b if they do so (perhaps capturing the role of sons in providing retirement income documented in Ebenstein and Leung, 2010; Ebenstein, 2014). They are also able to to engage in sex selection if they are willing to pay psychic cost γ_s capturing the disutility associated with sex selective abortion. Lifetime utility is

$$U = u(c) + n\gamma_n + b\gamma_b - s\gamma_s \tag{1}$$

[11] For the interested reader, there is a huge literature on the reforms. See e.g. Lardy (1983); Perkins (1988); Sicular (1988); Oi (1991); Huang (1998).

[12] Cotton and oilseeds are the two most important cash crops in China in the early 1980s. It is likely that output also grew faster in areas suited to other cash crops such as tea and sugarcane. However, data limitations prevent me from investigating this.

where b and s are dummy variables indicating the presence of one or more sons and whether the household has engaged in sex selection. I restrict γ_n, γ_b, $\gamma_s > 0$.

Both parents work. The father supplies one unit of labour inelastically at the prevailing male wage w. The mother divides one unit of labour between paid work and childrearing. Rearing each child is assumed to take $1/\psi$ units of time, with $\psi > 1$. The mother supplies $1 - \frac{n}{\psi}$ units of labour to the market at female wage w^*. There are no intrinsic costs to having children beyond the opportunity costs the mothers time. However, to capture the effect of the 1CP, a fine of $F \geq 0$ is imposed for each child after the first one.[13] The lifetime budget constraint is thus

$$c = w + \left(1 - \frac{n}{\psi}\right) w^* - \mathbb{1}(n > 0)(n-1)F. \tag{2}$$

Households maximise utility subject to this constraint. Unlike in the simplest models of fertility, here households have preferences over the gender of their children, and the ability to choose that gender. This means that fertility decisions should be made sequentially (as they are in life). Let the household have $N \leq \psi$ opportunities to have a child. In each period the household chooses whether to have a child and then, if they do so, whether to engage in sex selection. For simplicity, let the probability of having a boy in the absence of sex selection be $1/2$. Sex selection is assumed to work perfectly, and increases the probability that a boy is obtained to 1.[14]

The solution to the households fertility problem can be solved by backward induction. For ease of exposition, let $N = 2$; the household can have at most two children. The timing of the game is thus

1. *Households choose whether to have a child. If they choose to have a child, they also choose whether to select the sex of that child.*
    - *If they choose to have a child, and also choose its sex, they get a child of the sex of their choice.*
    - *If they choose to have a child, but didn't choose its sex, they get a girl with probability 1/2 and a boy with probability 1/2.*
    - *If they choose not to have a child, they get no child.*
2. *Households choose whether to have a child. If they choose to have a child, they also choose whether to select the sex of that child.*
    - ...
3. *Income earned, fines paid, and payoffs realised according to (1).*

A few things are immediate. First, given the costs and benefits of sex selection, households will only sex select in favour of a girl if the Sex Selection Condition holds

$$\gamma_s < \frac{\gamma_b}{2}. \tag{3}$$

[13] This approximates the situation faced by childless urban households, or Han rural households with one daughter for most of the history of the 1CP.

[14] Note, in this setup, sex selection does not entail an opportunity cost in terms of total fertility—it does not take time.

Second, as sex selection is costly, and HH's only gain additional utility from the first boy, HH's never sex select if they already have a boy. This is consistent with the patterns of sex ratios in figure 1. Third, because a first boy is more valuable than a second, a second child has a higher benefit if the first is a girl. Fourth, as sex selection is costly and households could get a boy by chance, sex selection is never optimal if the household will subsequently have another child. (The flip side of this is that, when the Sex Selection condition holds, households always sex select the gender of their last child.)

Consequently, whether the Sex Selection condition holds or not, when households are limited to two children there are four basic fertility strategies: (1) have no children; (2) have one child; (3) have two children if the first is a girl, otherwise have one child; and (4) have two children. If the Sex Selection condition holds, then households select the sex of their *last* child. Table 1 outlines each of the strategies and the associated payoffs. Which of the 4 strategies chosen depends on the parameter space.

The model has the following comparative statics for fertility. Higher male incomes, w, are associated with higher fertility. Higher female incomes, w^*, have an ambiguous effect on fertility: the income effect increases the demand for children, but the substitution effect increases their relative price. Given the budget constraint, the income effect is likely to dominate at relatively low levels of fertility, perhaps as in 1980s China, whereas the substitution effect dominates at high levels of fertility. This is consistent with the fertility transitions observed in the developed countries, where increases in women's incomes have been accompanied by reductions in total fertility (Becker, 1965; Willis, 1973; Heckman and Walker, 1990). A more severe One Child Policy, F, reduces fertility as one would expect, while higher utility from children, γ_n, and boys, γ_b, have opposite effect. Falls in the cost of sex selection encourage having a single child relative to other options. As relatively few households are childless, this suggests that the rollout of diagnostic ultrasound may have contributed to fertility decline during the 1980s.

The model has the following comparative statics for sex selection. If the cost of sex selection is high, then households never engage in sex selection. If it is low, then higher male incomes will tend to increase fertility. Whether this increases or decreases sex selection depends on whether most of the increase in fertility is due to the extensive margin effect of formerly childless couples deciding to have children (which increases sex selection) or the intensive margin effect of couple increasing their total fertility (which reduces it). Given that relatively few couples are voluntarily childless in China, it is likely the intensive margin effect dominates and sex selection falls. As with fertility, the effect of increases in female income are ambiguous. Falls in the cost of sex selection, γ_s, increase the propensity for sex selection (as shown in Chen et al., 2013; Lin et al., 2014). Consistent with the findings of Ebenstein (2010), the 1CP increases sex selection, the more so when fines F are high relative to income. Falls in preferences for children, γ_n, increase sex selection provided the main margin of adjustment is from two children to one, rather than one child to none. Declines in son preference unambiguously reduce sex selection.

The predictions of this simple model are broadly consistent with the findings of other papers considering the relationship between fertility, sex selection, the one-child policy and China's missing women. If the mechanism underpinning the model is correct, increases in incomes due to China's agricultural reforms are likely to have

increased fertility and reduced the demand for sex selection. These are the two main predictions I take to the data. I will also verify that parental education, a marker for income, decreases fertility much more in mothers than in fathers.[15]

4 Empirics

In this section, I first explore whether households more able to afford the fines associated with the 1CP had different fertility behaviour. Then, after showing that they do, I ask whether fertility transition had a causal effect on the increasingly unbalanced sex ratios observed in China.

The model above describes a world where low incomes are a constraint on fertility (as we might expect them to be in a world where there is substantial surplus rural labour, and the cost of having children is high). Because incomes are likely endogenous to preferences over fertility and sons, we cannot simply compare the fertility decisions of rich and poor. Instead, I obtain variation in incomes using a natural experiment provided by China's agricultural reform. In particular, I will use the fact that parts of China suited to growing cotton and oilseeds were better placed to benefit from the liberalisation of their planting than areas suited to grain and, as a result, had higher post-reform income growth.

The analysis will thus proceed in three steps. The first step is to confirm the findings of Marden (2015), and show that counties of China suited to cash crops did benefit more from China's agricultural reforms, and these gains were reflected in household income. Next, I test whether women residing in these counties increased their fertility relative to their counterparts elsewhere. Finally, I explore the gender composition of these births.

4.1 Incomes

The first step of my analysis verifies that counties endowed with land suited to cash crops did indeed benefit more from China's agricultural reforms. I estimate the following difference-in-differences specification, which compares post-reform income growth in areas more or less suited to cash crops

$$Y_{it} = \alpha_i + \gamma_t + \beta(Suitability_i \times Post_t) + \varepsilon_{it} \tag{4}$$

where Y_{it} is either net rural income or GDP per capita, α_i are county fixed effects, γ_t are province-time fixed effects, $Suitability_i$ is my normalised measure of suitability (the construction of which is described below), and $Post_t$ is a dummy for years after the reforms. These fixed effects control for all time invariant differences between counties, as well as differential province level economic performance.

The income data is from China's Provincial Anniversary Yearbooks. Although most provinces produced Anniversary Yearbooks, only a subset of them provided historical statistics at the county-level before and after the reforms. In all, I have comparable data for most of Gansu, Guizhou, Hebei, Jiangxi, Xinjiang and Zhejiang, and for some

[15] That it reduces fertility for males in contradiction to the predictions of the model, is likely due to the unmodelled relationship between preferences between cultural preferences for fertility and incomes which the variation in incomes below is intended to allow me to avoid.

prefectures Sichuan and Shanxi. My data on fertility and sex selection will span 1973-1990, so for incomes I use data for the years around this period for which fairly balanced data is available (1970, 1978, 1985, 1990). The reforms began in late 1978, and were largely completed by 1984, so I will treat 1970 and 1978 as 'pre-treatment' years, and 1985 and 1990 as 'post-treatment'.

To identify the parts of China suited to cash crops I use the measure of suitability developed in Marden (2015). This measure defines suitability at location ι as the ratio of the value of output in cash crops relative to that in grain

$$SCC_\iota = \frac{\max\{\Psi_{c\iota} p_c\}_{c \in \mathbb{C}}}{\max\{\Psi_{c\iota} p_c\}_{c \in \mathbb{G}}} \tag{5}$$

where p_c is the price of crop c, Ψ_c yield and $\mathbb{C}$ and $\mathbb{G}$ are the set of cash crops and grains respectively. Because the measure is relative, it captures the proportional gains from switching to cash crops—if the revenue productivity of cotton is twice that of wheat, switching to cotton should double output. I briefly describe the key inputs into the measure here, but for a fuller discussion of the merits of each input see Marden (2015).

The empirical implementation uses the measure uses above-quota prices from Sicular (1988) for the year 1978. Above-quota prices were the prices faced by farmers for deliveries to the government in excess of their mandated quotas—the marginal price faced by farmers. I use prices from 1978 which have the virtue of preceding the change in agricultural output which followed the reforms. I obtain productivities from the Food and Agriculture Organisation's Global Agro-Ecological Zones (GAEZ) database, which provides theoretical estimates of gross physical output per hectare under optimal growing conditions at a high spatial resolution.[16] The productivities are based on agronomic models which give measures of potential crop yield based on climatic conditions, soil type, elevation and gradient. The data provides productivities for a number of scenarios, allowing for various level of inputs (fertiliser use, mechanisation, modern seed varieties) and irrigation. As irrigation is widespread in China, I use the productivities based on 'intermediate inputs' and 'irrigation'.

Figure 2 is a map depicting land suitability to cash crops for the whole of China; significant variation exists both across and, crucially for my empirical strategy, within provinces. Because the finest geographic unit I can assign an individual to in the census is a county, I aggregate the cell level measure to the county level by taking the simple average of cell midpoints within the county $SCC_i = n_i^{-1} \sum_{\iota \in i} SCC_\iota$ where n_i is the number of fertile cell midpoints in county i. County boundaries for the 1990 census were obtained from University of Michigan China Data Center (2006). For ease of interpretation of estimated coefficients, I normalise this aggregated measure.

[16] In Beijing a cell represents an area about 6.5km square, cells are larger towards the equator, so in Shanghai a cell is approximately 8km square. Within economics, the use of the GAEZ data was pioneered by Nunn and Qian (2011), and the data has been used subsequently by several authors including Costinot and Donaldson (2014); Bustos et al. (2013); Marden (2015)

4.1.1 Results for rural incomes

Table 3 contains estimates of the relationship between my measure of suitability and post reform income growth at the county level. Both net rural income per capita and GDP per capita grew faster in counties best placed to benefit from planting cash crops (Columns 1 and 3). A 1 standard deviation increase in suitability for cash crops is associated with an additional 20% increase in GDP per capita after the reforms, and an additional 11% increase in net rural income per capita. For comparison, for the average county in my sample, real GDP per capita more than doubled between 1978 and 1990, while net rural income more than tripled.[17] The estimated coefficients are significant at the 1% and 5% level respectively.[18]

My baseline specification includes county and province-by-time fixed effects to allow for differential trends by province and sweep out time invariant differences between counties. If I omit the county fixed effects, I can include suitability for cash crops directly, and check whether suitability was correlated with incomes prior to the reforms. Columns 2 and 4 provide these results: suitability for cash crops was uncorrelated with incomes before the reforms. While there were no differences in incomes pre-reform, places suited to cash crops were considerably richer than places suited to grains in the years after the reforms. (The coefficients describing the post reform gains from specialising in the regressions without individual fixed effects are very similar to those obtained from specifications including fixed effects, but are estimated with somewhat less precision.)

The lack of a pre-existing correlation means that prior to the reforms, areas more or less suited to cash crops were comparable in terms of economic prosperity. This is supportive of the idea that suitability for cash crops was distributed throughout China in a manner as good as random. In a similar vein, Marden (2015) contains a large number of robustness checks for similar empirical results including, confirmation that counties more or less suited to cash crops were following parallel trends in a wide range of economic outcomes in the years prior to the reforms; showing that counties suited to cash crops actually started to grow more cash crops in the wake of the reforms; and showing that the results are robust to alternative ways of calculating suitability for cash crops. For brevity, I do not replicate these types of results here.

4.2 Fertility

The second step is to establish whether higher incomes affected household fertility. I estimate the effect of incomes on fertility using a difference-in-difference strategy exploiting only within woman variation. My baseline estimating equation is the following linear probability model

$$Birthevent_{it} = \alpha_i + \gamma_{jt} + \phi Birthevent_{it-1} + \beta \left(SCC_i \times Post_t \right) + x_{it}^{'}\theta + \epsilon_{it} \qquad (6)$$

[17] GDP per capita deflated by the national GDP deflator, net rural incomes by national RPI. Deflator and RPI obtained from the university of Michigan's China Data Center.

[18] Standard errors are two way clustered (Cameron et al., 2011) at the prefecture (the administrative division above the county) and the province-by-time levels to allow for autocorrelation of errors over time and space.

where α_i is a woman specific FE, γ_{jt} is a province-by-time FE allowing for differential trends in fertility by province, SCC_i is the suitability for cash crops of the county the woman resides in and $Post_t$ is a dummy taking a value of 1 for years after 1981.[19] I allow for a vector of time varying controls, x_{it}, including a full set of province-by-age fixed effects and controls for family composition. In my baseline specification, errors are clustered at the province level, allowing for auto-correlated errors over time and space. I create the fertility panel using the set of married women in the 1% sample of the 1990 population census. I focus on women who are either the head of household, or the spouse of the head of household, and who were born between 1945 and 1959. These are the women who I am able to reconstruct a full fertility history for. Table 2 contains summary statistics for this data.

One potential issue is that higher incomes could be related to fertility other than through an increase in the demand for children. Particularly, higher incomes could be associated with earlier marriage, which—given the rarity of childrearing out of wedlock in China—would lead to an increase in fertility in the earlier years of life, without, potentially, resulting in higher total fertility. To the extent that this earlier marriage was not driven by higher desired fertility, my estimates will be upward biased. The data in the 1990 census do not indicate that this is likely to be a problem: young women living in counties suited to cash crops are no more likely to be married than those elsewhere and no more likely to have had children, furthermore, the age of young women who are married or have children is uncorrelated with suitability for cash crops.[20] Nevertheless, I also provide results where the sample is restricted to women who are already married.

4.2.1 Results

Table 4 column 1 provides my baseline estimate of the effect of higher incomes on fertility. A one standard deviation increase in suitability for cash crops—my proxy for income—is associated with a post-reform relative increases in the likelihood of a birth of 3.2 percentage points per year. This estimate is significant at the 1% level.[21] The effect size is large, at 28% of the average annual fertility rate in my sample, but equates to a much smaller overall increase in fertility due to the observed negative impact of previous fertility on the likelihood of additional births. Consistent with Chinese household's preference for sons, the marginal reduction in fertility after having a boy is greater than that for a girl. Column 2 restricts the sample to women who have already had at least one child, and so captures purely the effect on fertility on the intensive margin. Because the sample includes women who have not completed childrearing, the estimate is a lower bound estimate of the intensive margin effect. Nevertheless, the

[19] The choice of 'treated' and 'non-treated' years is not crucial for my analysis. The reforms were implemented gradually between 1978 and 1984 and for decommunalisation at least, the bulk of the changes took place between 1981-83 (Lin, 1992). In light of this 1982 provides a reasonable first 'post-reform' year. Online Appendix table A2 provides supplementary results based on post-reform being defined as beginning in 1979 or 1985 with similar results.

[20] See Online Appendix table A1 for these results and Online Appendix section B.1 for an extended discussion.

[21] In my baseline specification, standard errors are clustered at the province level. Because the number of provinces (29) is slightly below the number desirable to be reasonably sure of consistency, I also provide my baseline results with errors clustered at the prefecture level (Column 5). The standard errors based on clustering at the prefectural level are significantly smaller.

estimated coefficient indicates that at least 2/3 of the effect is due to increases in fertility along the intensive margin and, despite being a lower bound, the intensive margin effect is not statistically significantly different from the full effect estimated in column 1.

Columns 3 and 4 provide estimates for male and female births. The increase in total fertility is split almost equally, with a one standard deviation increase in suitability associated with a 1.7 or 1.6 percentage point increase in the probability of a male or female birth respectively—additional births are split between male and female births at the natural rate. Given China's imbalance in births, these additional births represent a reduction in gender bias. Consistent with gender discrimination, and the theoretical framework provided, the effect of past births effect the likelihood of new male and female children quite differently: more boys primarily reduce the probability of male births, more girls only reduce the probability of more female births.

As with any difference-in-difference strategy, the identification assumption is that of parallel trends. That is, that women living in counties more or less suited to cash crops, would have had identical changes in fertility in the absence of the reforms. While it is not possible to verify this assumption directly, we can use the pre-reform years as a placebo test of the parallel trends assumption. To implement this, I estimate a version of Equation 6, estimating a coefficient on suitability for each year in the data

$$Birthevent_{it} = \alpha_i + \gamma_{jt} + \phi Birthevent_{it-1} + \sum_{s \neq 1973} \beta_s \left(SCC_i \times I\{Year == s\} \right) + \epsilon_{it} \quad (7)$$

as the reforms began in late 1978, we would expect $\beta_t = 0 \; \forall \, t \leq 1978$. Any effects of the reform ought to appear gradually after 1978, with the full effect potentially not observed until a few years after the completion of the agricultural reforms in 1984.

Figure 3 plots these estimated coefficients and associated 95% confidence intervals for each year in the data. Prior to the reforms there were no differential trends in fertility, whereas after the reforms areas suited to cash crops experienced a substantial increase in their relative fertility. This is consistent with the parallel trends assumption.

The framework provided in Section 3 suggested that higher incomes for women would likely reduce fertility, whereas higher incomes for men would probably increase it. The census data do not include a measure of household or individual income, however I can observe parental education, which is likely to be closely related to income. Unfortunately, even more than income itself, education is likely to be a strong correlate of weaker cultural preferences for children. With this in mind, given that the opportunity cost of having children in terms of lost income is much higher for women than for men, we should expect the relationship between education and fertility to be *more negative* for women than for men. Table 5 provides these results. High education for either parent negatively predicts fertility. (The measure of 'high education' is attendance at junior middle school or higher, a level achieved by just over a third of women and just under half of men.) However, when included together, maternal education is the principle driver of this lower fertility. This is consistent with the opportunity cost theory of fertility.[22]

[22]Because parental education is time invariant in my data, I can no longer include woman specific fixed effects. I instead include a battery of county, parental registration and ethnicity fixed effects. This change in specification attenuates the coefficient on my main variable of interest but it remains substantial and statistically significant at the 1% level.

Households in parts of China that were able to benefit from growing cash crops, had higher incomes and higher fertilities in the wake of the reforms. However, economic growth is more usually associated with lower fertility. It is hence worth considering why in this case, higher incomes led to higher fertility.

One possibility, discussed in more detail in section 2, is that higher incomes reduced the effectiveness of the One Child Policy (1CP). Because fines associated with breaching the 1CP were primarily set at the province level, then within a province—the level the empirical analysis is at—households with higher incomes will be more able to afford these fines. Alternatively, because local authorities had some discretion over how draconian enforcement of the 1CP was, the governments of more successful areas may have been able to be more liberal. In either case, higher incomes could have effectively reduced the cost of breaching the 1CP.

I test this proposition by exploiting the fact that the 1CP did not affect all populations equally. In particular, non-Han households were subject to more liberal policies than their Han counterparts. If the increase in fertility was driven by an interaction between higher incomes and the 1CP, we would thus expect Han households to be more strongly affected than non-Han Households. Table 6 provides this test. Column 1 restates my baseline results. Column 2 restricts the sample to Han mothers. Column 3 restricts the sample to non-Han mothers. Fertility increases strongly in Han households and the estimated coefficient is significant at the 1% level. In non-Han households, the effect on fertility is one-third of the size and statistically insignificant. These results are consistent with higher incomes alleviating constraints on fertility due to the 1CP, particularly for households that are the most constrained.

These results also mitigate against another potentially important alternative explanation for fertility change: that producing cash crops may have changed the calculus in the quantity-quality tradeoff (e.g. Becker and Lewis, 1974). For this to be a concern, shifting to the production of cash crops ought to have encouraged quantity rather than quality. The classic argument suggests that as the return to skill increases, parents choose to have fewer higher quality children. Thus, if growing cotton were less skill intensive than grain, the calculus could move in favour of quantity. This does not appear to be the case: the distribution of education of workers in cotton is almost identical to that of grain (figure 4), with if anything, workers in cotton being slightly more educated.[23] (The intuition provided by this visual inspection is confirmed by an unreported ordered probit regression which indicates that workers in cotton marginally better educated.) As Marden (2015) documents a substantial subsequent increase in non-agricultural output in these areas—activity that is relatively human capital intense—it seems safe to say that, if anything, the returns to human capital is likely to have increased. We might also expect to see the quantity-quality calculus shift towards quantity if the demand for child labour increased, however, perhaps because of the surplus of adult labour in China, children do not form a major part of the rural labour force in China (Bowlus and Sicular, 2003).

[23] Unfortunately the data do not permit the same comparison with respect to oilseeds.

4.3 Sex selection

The level of sex selection increased substantially in 1980s China. This paper asks whether fertility transition was an instrumental factor in this change. We have seen in the previous section that, within a province, higher incomes relaxed the constraints on fertility due to the One Child Policy and led to higher birth rates. I now ask whether this reduced the level of sex selection.

As shown in figure 1, the increasing share of boys was entirely due to changes in birth shares for higher order births in households without a son. I exploit this differential discrimination to allow the use of a triple difference strategy. The first difference is in the timing of the reforms. The second is in geographic suitability for cash crops (my proxy for higher incomes/fertility. The third is the presence of older brothers.[24] To implement the third difference, I restrict my attention to second-born children, as sex discrimination against boys for some third born children would complicate the interpretation of the estimated coefficients.

In practice, my empirical implementation of the triple difference strategy utilises a large set of fixed effects. The estimating equation is

$$Girl_{ijkt} = \alpha_{jk} + \gamma_{jt} + \delta_{kt} + \beta(SCC_i \times Post_t \times NoBrother_i) + \epsilon_{ijkt} \tag{8}$$

where α_{jk} are county-by-no-brother fixed effects, γ_{jt} are county-by-birth-year fixed effects, and δ_{kt} are province-by-birth-year-by-no brothers fixed effects. The outcome variable $Girl_{ijkt}$ is a dummy taking a value of 1 if the child is a girl.

The inclusion of county-by-no-brother fixed effects allows for time invariant difference in son preference and sex selection at the county level. The county-by-birth-year fixed effects allow for county-year specific shocks to sex ratios including economic and environmental shocks in-utero and through childhood, which differentially effect the survival of boys and girls. For instance, this allows for the fact that male foetuses are more fragile and thus less likely to be born in times of hardship (Andersson and Bergström, 1998). The province-by-birth-year-by-no-brother fixed effects flexibly allow for differential changes in son preference and sex selection by province. (This last set of fixed effects are the triple difference analogue of province-by-time FE's in a standard difference-in-difference specification.)

I impute the required data from the 1% sample of the 1990 population census. Because my analysis is of a surviving cohort in 1990, rather than actual births, it is important to highlight that any effect on the gender uncovered captures both sex selection at birth *and* discrimination between birth and 1990. While my data do not allow me to estimate the effect of only antenatal sex selection, Anderson and Ray (2010) has shown that while China has hugely skewed sex ratios at birth, the level of discriminatory mortality during childhood is low. In light of this, my empirical results are likely to be driven by discrimination before birth, and throughout the paper I use the term 'sex selection' as a synonym for total discrimination.

The individuals of interest are those born in the years before and after the reform who have one older sibling. I recover this information from the presence of other children of household heads in the same household. This approach necessitates a

[24] The empirical strategy is closely related to that in Almond et al. (2013), although the specification and source of variation is quite different.

number of restrictions on the set of births used. First, I can only use households with a unique household head and spouse, which precludes the possibility of using births to households in communal living arrangements. Second, because in the census HH structure is defined relative to the household-head, I restrict myself to the children of HH heads: there are no grandchildren or cousins of HH heads in the sample. Third, I exclude children in households which have migrated in the five years prior to the census as I don't know where they migrated from, and hence, am unable to determine how much they benefited from the reforms prior to migration.[25]

It has sometimes been argued that China's demographic data undercount the youngest children and, in particular, young girls, with the consequence that sex ratios at birth look worse than they really are. Banister (2004, p. 20) reviews the evidence for this proposition and concludes that 'in the 1990 census [the primary source of demographic data employed in this paper], there was no more undercounting of young children of one sex than the other'. Given this, and the empirical strategy employed, any undetected undercounting is unlikely to be problematic.[26]

4.3.1 Results

Table 7 column 1 provides my baseline triple-difference result. Counties with more rapid post-reform increases in incomes and fertility, experienced relative falls in gender discrimination at birth after the reforms. A one-standard deviation increase in suitability for cash crops is associated with a 2.4 percentage point increase in the probability of a second born child being a girl if it has an older sister (when discrimination commonly occurs) compared to the probability of having a girl when the household already has a boy (and sex selection is extremely rare). This estimate is significant at the 1% level and should be compared to an overall deficit in the number of girls of between 8 and 10 percentage points.[27]

Columns 2 and 3 introduce controls for maternal and paternal age, education and industry, by including a full set of dummies for each. The estimated coefficients on my variable of interest are completely unchanged.

Because birth order is imputed from a child's age relative to that of other children in the household, rather than from information directly provided, it may be imputed with error where one or more children have died or are temporarily or permanently out of the household. I can greatly reduce problems of this type by restricting the sample to children born to households where the number of children present is exactly equal to the number the mother reports ever having been born. Whether this is desirable or not depends on when the child went missing or died. For the younger siblings of children who died in infanthood, with respect to the 1CP it is as if the second born child is the

[25] Unfortunately, the census does not include complete information on migration history, so some households may have migrated before 1985. However, migration was very low during this period so the resultant bias should be minimal.

[26] Even, if there were some undercounting of girls, this would only be a problem for my empirical results if the undercounting was (a) related to whether the firstborn child was male or female, and (b) correlated with my proxy for post-reform income growth.

[27] In my baseline specification, standard errors are clustered at the province level. Because the number of provinces (29) is below the number desirable to be reasonably sure of consistency, in column 5 I provide my baseline results with errors clustered at the prefecture level. The standard errors are slightly smaller so I proceed with the more conservative clustering at the province level.

first born and the third born child the second born. Hence, in my baseline specification I choose not to impose this restriction. My results are nevertheless almost completely unchanged if children from households where the number of stated surviving children does not match the number of children in the household are excluded (column 4).

The identification assumption is that of conditional parallel trends. Within a province there would have been no differential change in the probability of a second child with no brothers being born (and surviving) in areas suited to cash crops (relative to those born in places suited to grain) beyond any changes in the sex ratio of those born in families which already had a son. This identification assumption does not require that higher incomes or growing cash crops had no direct effect on the probability that a second child is male, only that the direct effect is the same regardless of the gender of the first child. For instance, changes in sex ratios as a result of higher incomes, changing diets or exposure to pollution, pesticides or fertiliser are only problematic if they affect the gender of second born children conditional on the gender of the first born.

Although it is not possible to directly verify the conditional parallel trends assumption, the pre-reform years can act as a placebo test. To implement this, I estimate the triple-difference analogue of equation 7 which allows for a differential effect of suitability for each birth year

$$Girl_{ijkt} = \alpha_{jk} + \gamma_{jt} + \delta_{kt} + \sum_{s \neq 1973} \beta_s \left(Suited_i \times I\{BirthYear == s\} \times NoBrother_i \right) + \nu_{ijkt}. \tag{9}$$

As the reforms begin in late 1978, we would expect $\beta_t = 0 \; \forall \, t < 1978$. Then as the reforms are implemented between 1978 and 1984 we would expect any effects on sex ratio to begin to appear. However, to the extent that increased demand for sex selection only translated into action after the availability low cost sex detection technology, the effect on gender might not be apparent until later in the 1980s.

Figure 5 plots the triple-difference treatment effect of suitability on the sex of second born children for each year in the data. Prior to the reforms there was no differential trend in birth shares, whereas after the reforms a substantial differential effect becomes apparent. This is consistent with the key assumption of conditional parallel trends.

Interestingly, the difference in sex ratios only begins to emerge in the mid-1980s around the end of the reforms rather than during them. This is approximately contemporaneous with the widespread availability of low cost sex detection technology (Chen et al., 2013) and suggestive that the lower demand for sex selection generated by higher fertility only became an important determinant of gender ratios once this demand (or lack thereof) was relatively easy to act upon. However, because of this timing we must be concerned that my proxy for income and fertility growth was correlated with some latent variables affecting the demand for sex selection. We have already seen that incomes are essentially uncorrelated with suitability prior to the reform suggesting that these areas do not differ substantially economically. In table 2 I report correlations between suitability for cash crops and education, registration status, migration and minority status. Only minority status is correlated with any degree of strength: woman are more likely to be Han and less likely to be minorities. Because non-Han Chinese had larger families and were less affected by the 1CP, they are likely to have had weaker latent demand for sex selection, indicating that, if anything, my estimates are probably

downward biased.

Table 8 once again exploits the fact that fertility restrictions were stricter for Han Chinese than ethnic minorities to confirm that it the effect on sex ratios is coming through the same groups that also increased their fertilities. Column 1 restates my baseline results. Column 2 restricts the sample to Han. Column 3 restricts the sample to non-Han. Higher incomes (proxied by suitability for cash crops) are associated with more girls amongst Han Chinese, but not amongst ethnic minorities. This is consistent with higher rural incomes reducing gender bias only where higher incomes enabled an increase in fertility

4.3.2 Women's labour market outcomes

If women had a comparative advantage in growing and distributing cotton and oilseeds, specialisation in these crops may have affected their relative labour market prospects. While the differential labour market prospects do not appear to be an important factor in cross-national or cross ethnic differences in sex selection (Ebenstein, 2014), in early 1980s China increases in women's relative earnings do appear to reduce sex selection (Qian, 2008).[28]

If growing cotton and oilseeds improved women's labour market outcomes, then this would provide an additional explanation for the observed decline in sex selection (but not for the increase in fertility). One empirical measure of comparative advantage, is the extent to which women are represented in that industry relative to others. On this basis, women appear to have a modest comparative advantage in cotton farming, as they make up 57% of cotton farmers compared to 47% of all farmers.[29] (Unfortunately, the data do not include oilseed farming as an occupation.) However, it is not clear whether this differential translates into improved labour market outcomes.

One labour market outcome I can observe is labour force participation. If women's labour market prospects had improved, we might expect higher relative levels of female labour force participation in areas suited to cash crops. This is not apparent in the data. Table 9 contains cross sectional regressions of labour force participation on suitability for cash crops of an individuals county of residence, a female dummy and an interaction term based on the 1990 census. Labour market participation is only weakly correlated with female labour market participation in 1990—a one standard deviation increase in suitability as associated with a 1 percentage point increase in participation—and this association is not statistically significant. Adding province specific fixed effects, and fixed effects for household size, education and registration status do not materially affect the estimated correlation.

Of course, if counties suited to cash crops had low female labour market participation prior to the reforms, the cross section could be hiding a differential change in the economic status of women. We have already seen that rural incomes and GDP per capita were uncorrellated with suitability prior to the reform, so we would perhaps not expect labour participation to be correlated either. Nevertheless, we can check

[28] One explanation for this apparent inconsistency is that in general improved labour market prospects for women are likely to capture two offsetting effects: a decline in fertility due to the higher opportunity cost of having children which we would expect to increase sex selection and an improvement in the status of women which could reduce it.

[29] In the 1% sample of the the 1990 census. However there is an almost identical difference in 1982.

whether there were differential changes in participation by aggregating the census microdata from 1982 and 1990 to the prefectural level.[30] Table 10 contain difference-in-difference estimates of the effect of suitability for cash crops of female labour force participation based on this data. There is no statistically significant effect on female labour force participation and the estimated coefficients are small and of inconsistent sign.[31] In Column 1, we see there there were no pre-existing differences in labour force participation. Columns 2 and 3 add richer sets of fixed effects. Columns 4 and 5, allow for triple differences with respect to male labour force participation.

Thus, although the data suggests that women may be somewhat overrepresented in cotton farming, this does not appear to have translated into real differences in labour market outcomes. The role of differential income shocks by gender is hence likely limited. The previous results for fertility are consistent with this. Improved opportunities for women have generally been associated with declining fertility and, consistent with this, we have seen the status of women is correlated with lower fertility in section 4.2. Despite this, the higher incomes in this paper are associated with higher fertility. Although it is not possible to decisively rule out *any* role for gender biased technical change, the pattern of results are much more consistent with a direct link from fertility to sex selection.

5 Conclusion

Parts of China more suited to cash crops had faster income growth in the wake of China's agricultural reforms. Households in these areas increased their fertility, most likely due to higher incomes effectively relaxing the One Child Policy. Higher desired fertility meant that households were more likely to have a boy by chance and, as a consequence, fewer households felt the need to select the sex of their offspring in order to satisfy their cultural preference for sons.

Interestingly, if we compare the timing of the results, fertility increases almost immediately, but the fall in sex selection mirrors the timing of the rise in sex selection in China as a whole (in that it becomes apparent from the mid 1980s). This is consistent with the timing of the increased availability of pre-natal sex selection documented by Chen et al. (2013). Thus while fertility decline may have driven the demand for sex selection, this latent demand may not have been fulfilled until the price of sex selection became sufficiently low.

The increasingly biased sex ratios observed in Asia thus appear to have been due to the intersection of two trends: fertility transition and the availability of ultrasound. This explains why the fall in fertility in 1970s China was not accompanied by a dramatic change in sex ratios, but also suggests that this fall in fertility may have been responsible for an increase in latent demand for sex selection that only manifested once the price of sex selection had fallen sufficiently.

[30] Unfortunately, county level identifiers are not included in the 1982 census micro data (and male and female labour force participation are not provided separately in the county level summary data) so the prefecture is the finest geographic partition available. The construction of this data is discussed in Online Appendix A.

[31] Because 1982 is somewhat after the start of the agricultural reforms—although, for at least the decommunalisation of agriculture, the bulk of the liberalisation too place between 1982 and 1984—these estimates may be biased towards zero. Given the inconsistent signs this is not too much of a concern.

References

Abrevaya, J. et al. (2009). Are there missing girls in the united states? evidence from birth data. *American Economic Journal: Applied Economics*, 1(2):1–34.

Almond, D. and Edlund, L. (2008). Son-biased sex ratios in the 2000 United States census. *Proceedings of the National Academy of Sciences*, 105(15):5681–5682.

Almond, D., Li, H., and Zhang, S. (2013). Land reform and sex selection in China.

Anderson, S. and Ray, D. (2010). Missing women: age and disease. *The Review of Economic Studies*, 77(4):1262–1300.

Andersson, R. and Bergström, S. (1998). Is maternal malnutrition associated with a low sex ratio at birth? *Human Biology*, pages 1101–1106.

Anukriti, S. (2014). The fertility-sex ratio trade-off: Unintended consequences of financial incentives.

Banister, J. (2004). Shortage of girls in China today. *Journal of Population Research*, 21(1):19–45.

Becker, G. S. (1960). An economic analysis of fertility. In *Demographic and economic change in developed countries*, pages 209–240. Columbia University Press.

Becker, G. S. (1965). A theory of the allocation of time. *The economic journal*, pages 493–517.

Becker, G. S. and Lewis, H. G. (1974). Interaction between quantity and quality of children. In *Economics of the family: Marriage, children, and human capital*, pages 81–90. UMI.

Bowlus, A. J. and Sicular, T. (2003). Moving toward markets? Labor allocation in rural China. *Journal of Development Economics*, 71(2):561–583.

Bulte, E., Heerink, N., and Zhang, X. (2011). China's one-child policy and 'the mystery of missing women': Ethnic minorities and male-biased sex ratios. *Oxford Bulletin of Economics and Statistics*, 73(1):21–39.

Bustos, P., Caprettini, B., and Ponticelli, J. (2013). Agricultural productivity and structural transformation. evidence from Brazil. *Mimeo*.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2).

Chang, C.-F., Lee, C.-F., McKibben, S. L., Poston, D. L., and Walther, C. S. (2005). *Fertility, family planning and population policy in China*. Routledge.

Chen, Y., Li, H., and Meng, L. (2013). Prenatal sex selection and missing girls in China: Evidence from the diffusion of diagnostic ultrasound. *Journal of Human Resources*, 48(1):36–70.

Costinot, A. and Donaldson, D. (2014). How large are the gains from economic integration? theory and evidence from US agriculture, 1880-2002. *Mimeo*.

Das Gupta, M. and Mari Bhat, P. (1997). Fertility decline and increased manifestation of sex bias in India. *Population studies*, 51(3):307–315.

Das Gupta, M. and Shuzhuo, L. (1999). Gender bias in China, South Korea and India 1920–1990: Effects of war, famine and fertility decline. *Development and Change*, 30(3):619–652.

Dubuc, S. and Coleman, D. (2007). An increase in the sex ratio of births to India-born mothers in England and Wales: Evidence for sex-selective abortion. *Population and Development Review*, 33(2):383–400.

Ebenstein, A. (2010). The missing girls of China and the unintended consequences of the one child policy. *Journal of Human Resources*, 45(1):87–115.

Ebenstein, A. (2014). Patrilocality and missing women. *Available at SSRN 2422090*.

Ebenstein, A. and Leung, S. (2010). Son preference and access to social insurance: evidence from China's rural pension program. *Population and Development Review*, 36(1):47–70.

Ebenstein, A. Y. (2007). Fertility choices and sex selection in Asia: Analysis and policy. *Available at SSRN 965551*.

Heckman, J. J. and Walker, J. R. (1990). The relationship between wages and income and the timing and spacing of births: evidence from Swedish longitudinal data. *Econometrica: journal of the Econometric Society*, pages 1411–1441.

Huang, Y. (1998). *Agricultural Reform in China*. Cambridge University Press.

Jayachandran, S. (2014). Fertility decline and missing women.

Lardy, N. R. (1983). *Agriculture in China's modern economic development*. Cambridge University Press.

Lin, J. Y. (1992). Rural reforms and agricultural growth in China. *The American Economic Review*, pages 34–51.

Lin, M.-J., Qian, N., and Liu, J.-T. (2014). More women missing, fewer girls dying: The impact of abortion on sex ratios at birth and excess female mortality in and excess female mortality in Taiwan. *Journal of the European Economic Association*, 12.

Marden, S. (2015). The agricultural roots of industrial development: forward linkages in reform era China. *Mimeo*.

McMillan, J., Whalley, J., and Zhu, L. (1989). The impact of China's economic reforms on agricultural productivity growth. *The Journal of Political Economy*, pages 781–807.

Nunn, N. and Qian, N. (2011). The potato's contribution to population and urbanization: Evidence from a historical experiment. *The Quarterly Journal of Economics*, 126(2):593–650.

Oi, J. C. (1991). *State and peasant in contemporary China: The political economy of village government*, volume 30. University of California Press.

Park, C. B. and Cho, N.-H. (1995). Consequences of son preference in a low-fertility society: imbalance of the sex ratio at birth in Korea. *Population and Development Review*, pages 59–84.

Perkins, D. H. (1988). Reforming China's economic system. *Journal of Economic Literature*, 26(2):601–645.

Qian, N. (2008). Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance. *The Quarterly Journal of Economics*, 123(3):1251–1285.

Rosenzweig, M. R. and Schultz, T. P. (1982). Market opportunities, genetic endowments, and intrafamily resource distribution: Child survival in rural india. *The American Economic Review*, pages 803–815.

Scharping, T. (2013). *Birth Control in China 1949-2000: Population policy and demographic development*. Routledge.

Sen, A. (1990). More than 100 million women are missing. *The New York Review of Books*.

Sen, A. (1992). Missing women. *BMJ: British Medical Journal*, 304(6827):587.

Sicular, T. (1988). Agricultural planning and pricing in the post-Mao period. *The China Quarterly*, 116(1):671–705.

Taylor, J. R. (1988). Rural employment trends and the legacy of surplus labour, 1978-86. *The China Quarterly*, 116:736–66.

University of Michigan China Data Center (2006). Historical China county population census data with GIS maps.

White, T. (2006). *China's longest campaign: birth planning in the People's Republic, 1949-2005*. Cornell University Press.

Willis, R. J. (1973). A new approach to the economic theory of fertility behavior. *The Journal of Political Economy*, pages S14–S64.

Yi, Z., Ping, T., Baochang, G., Yi, X., Bohua, L., and Yongpiing, L. (1993). Causes and implications of the recent increase in the reported sex ratio at birth in China. *Population and Development Review*, pages 283–302.

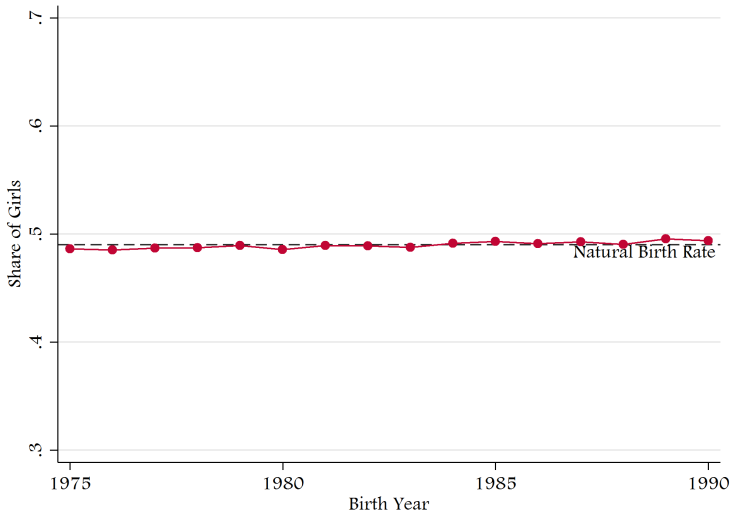

(1a) First Born Children

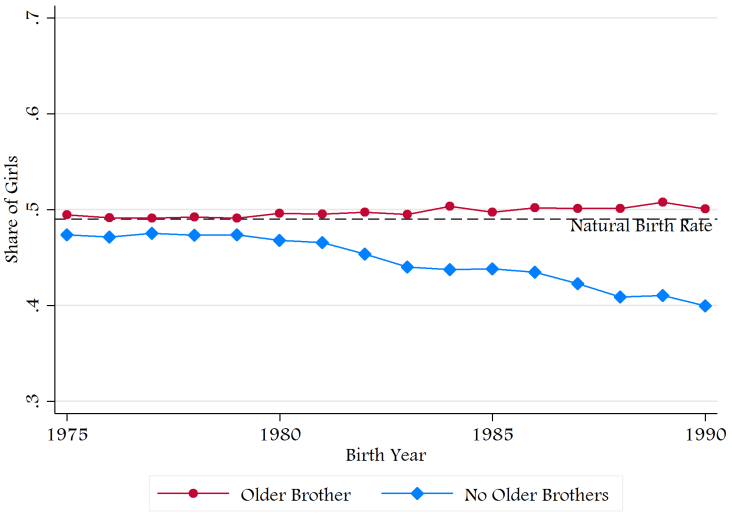

(1b) Second Born Children

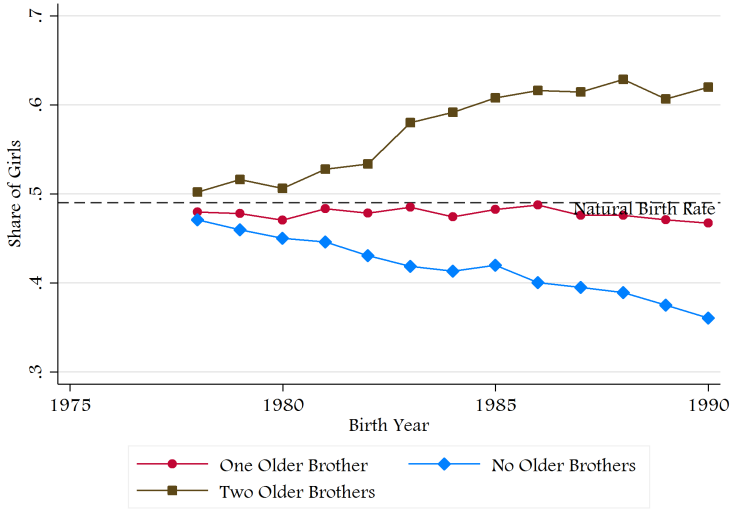

(1c) Third Born Children

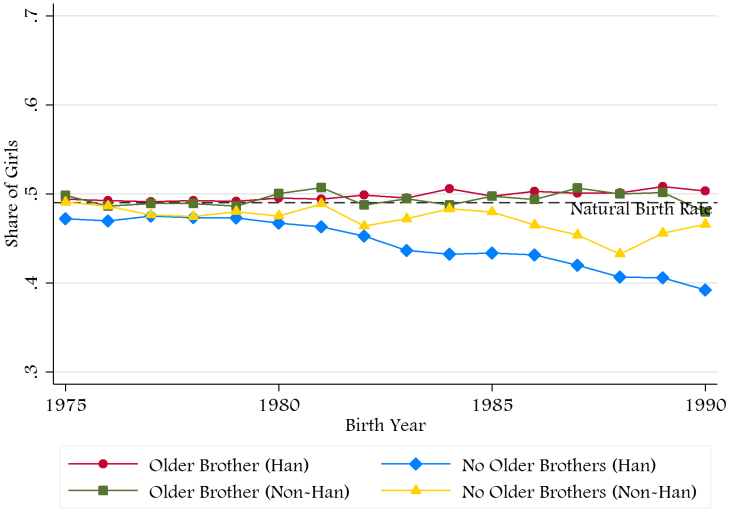

(1d) Second Born (Han vs. non-Han)

Figure 1: Share of Girls in Birth Year Cohorts Conditional on Birth Order, Family Composition and Ethnicity. (Source: Calculated from the 1% Sample of the 1990 Population Census)

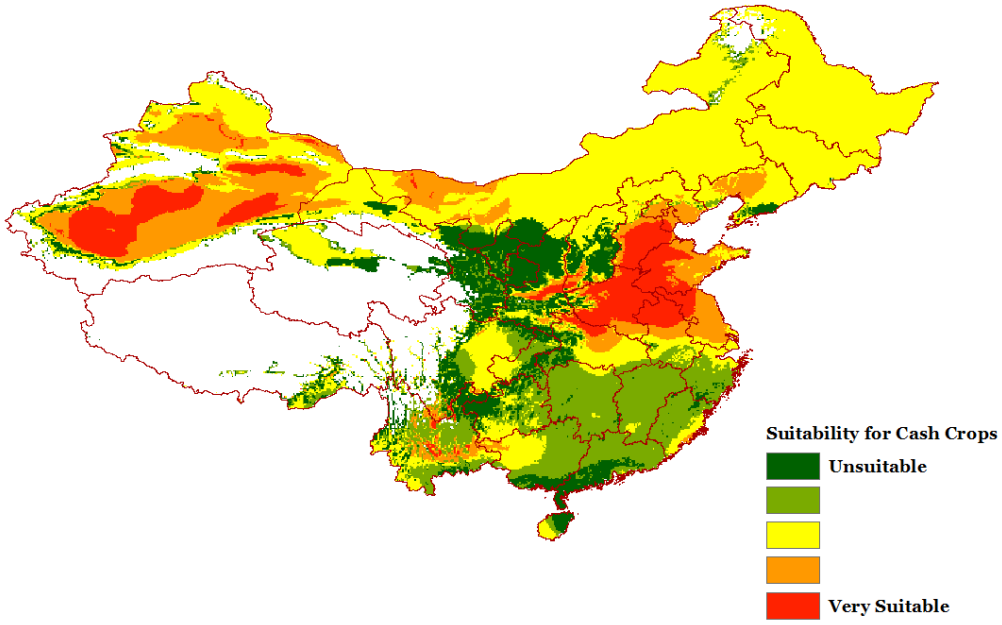

Figure 2: Suitability for Cash Crops

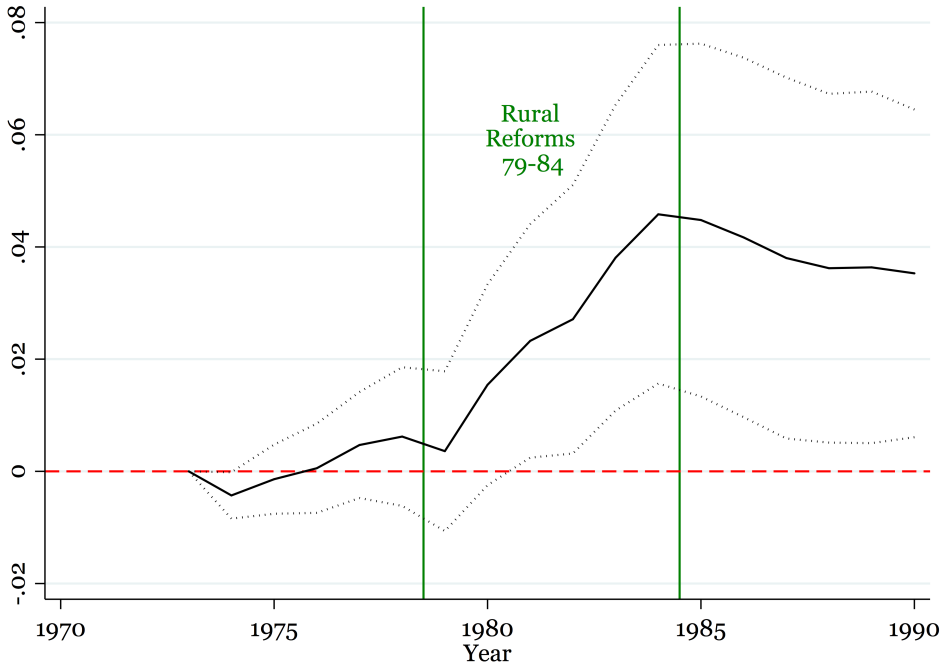

Figure 3: Effect of suitability for cash crops on fertility 1973-1990 and associated 95% confidence intervals.

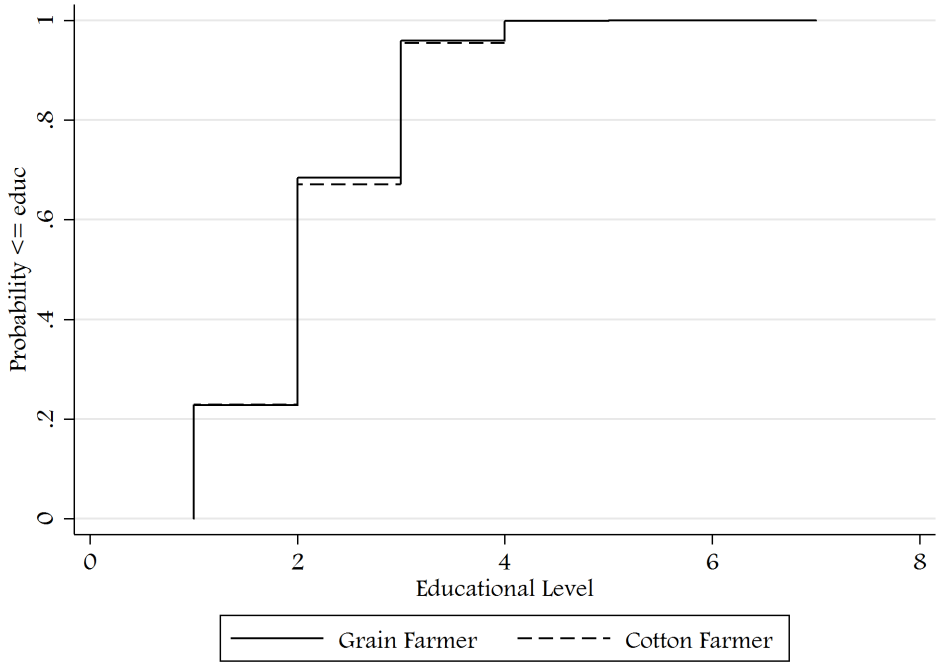

Figure 4: CDF's of Educational Levels for Grain and Cotton Farmers (0 "N/A" 1 "Illiterate" 2 "Primary" 3 "Junior Middle" 4 "Senior Middle" 5 "Technical School" 6 "Junior College" 7 "University")

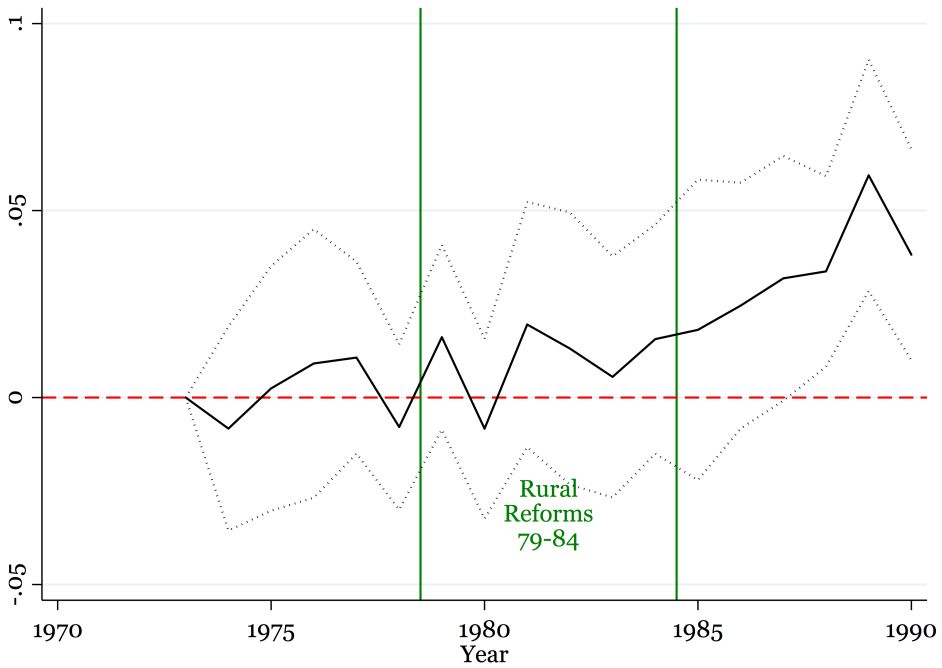

Figure 5: Triple difference estimates of the effect of suitability to cash crops on birth year 1973-1990 and associated 95% confidence intervals.

Table 1: Fertility Strategies

Strategy	Payoff
High Cost of Sex Selection $\left(\gamma_s > \frac{\gamma_b}{2}\right)$	
H1: Have no children	$u(w + w^*)$
H2: Have one child	$u(w + (1 - \frac{1}{\Psi})w^*) + \gamma_n + \frac{1}{2}\gamma_b$
H3: Have a child, if it's a girl have another one	$\frac{1}{2}\left[u(w + (1 - \frac{1}{\Psi})w^*) + u(w + (1 - \frac{2}{\Psi})w^* - F)\right] + \frac{3}{2}\gamma_n + \frac{3}{4}\gamma_b$
H4: Have two children	$u(w + (1 - \frac{2}{\Psi})w^* - F) + 2\gamma_n + \frac{3}{4}\gamma_b$
Low Cost of Sex Selection $\left(\gamma_s \leq \frac{\gamma_b}{2}\right)$	
L1: Have no children	$u(w + w^*)$
L2: Have one child and select a boy	$u(w + (1 - \frac{1}{\Psi})w^*) + \gamma_n + \gamma_b - \gamma_s$
L3: Have a child, if it's a girl have another one and select a boy	$\frac{1}{2}\left[u(w + (1 - \frac{1}{\Psi})w^*) + u(w + (1 - \frac{2}{\Psi})w^* - F)\right] + \frac{3}{2}\gamma_n + \gamma_b - \frac{1}{2}\gamma_s$
L4: Have two children, if the first is a girl select a boy for the second	$u(w + (1 - \frac{2}{\Psi})w^* - F) + 2\gamma_n + \gamma_b - \frac{1}{2}\gamma_s$

Table 2: Summary Statistics

x	$\bar{x}$	σ_x	N	$\rho_{x,relval}$
A. Annual Fertility Data (For Section 4.2)				
Any Birth	0.111	0.31	13.7m	
- male only	0.058	0.23	13.7m	
- female only	0.052	0.22	13.7m	
- conditional on 1 or more children already	0.096	0.30	9.3m	
- for non-Han mothers	0.138	0.35	795284	
Total Births by Age 40	2.762	1.11	136614	
B. Characteristics of Women in 1990				
At Least Middle School Education	0.362	0.48	3.0m	0.049
Rural Household Registration	0.778	0.42	3.0m	0.003
Han	0.930	0.26	3.0m	0.127
Migrated in Last Five Years	0.036	0.19	3.0m	0.012

Panel A based on fertility history of married women born 1945-1959 as described in Section 4.2 (all variables are dummies except total births). Panel B based on all women aged 21-65 in the 1% sample of the 1990 population census (all variables are dummies).

Table 3: Rural Incomes

	Ln Rur. Ypc		*Ln GDPpc*	
	(1)	(2)	(3)	(4)
Suitability for Cash Crops × Post 78	0.186***	0.165*	0.100**	0.116
	(0.046)	(0.094)	(0.050)	(0.084)
Suitability for Cash Crops		-0.050		0.008
		(0.074)		(0.057)
Observations	1871	1871	2214	2214
Counties	526	526	611	611
County FE	Yes		Yes	
Province x Year FE	Yes	Yes	Yes	Yes

Robust standard errors two-way clustered (Cameron et al., 2011) at the county and province-by-time level to allow for autocorrelation over time and space (*** p<0.01, ** p<0.05, * p<0.1). County level data on Net Rural Income Per Capita and Primary GDP obtained from 50th and 60th provincial anniversary yearbooks. Data covers the years 1970, 1978, 1985 and 1990 and counties in Gansu, Guizhou, Hebei, Jiangxi, Jiangsu (not Primary GDP), Xinjiang, Zhejiang (not Rural Income) and parts of Shanxi and Sichuan. Suitability for cash crops is the (standardised, county or province average) ratio of the value of output of the best cash crop to the value of output of the best grain crop. Marden (2015) explores the link between agricultural and industrial development in China and provides a large number of additional results indicating the robustness of suitability for cash crops as a predictor of income growth.

Table 4: Effects on Fertility (Baseline Results)

	Dummy for Birth of...				
	Any		Boy	Girl	Any
	(1)	(2)	(3)	(4)	(5)
Suitability for Cash Crops × 1982-1990	0.033***	0.020**	0.017***	0.016***	0.033***
	(0.011)	(0.009)	(0.006)	(0.005)	(0.005)
Had Child The Previous Year Dummy	-0.078***	-0.137***	-0.044***	-0.034***	-0.078***
	(0.007)	(0.010)	(0.004)	(0.003)	(0.003)
Number of Boys	-0.247***	-0.295***	-0.207***	-0.041***	-0.247***
	(0.008)	(0.007)	(0.005)	(0.004)	(0.003)
Number of Girls	-0.182***	-0.227***	-0.001	-0.182***	-0.182***
	(0.007)	(0.007)	(0.003)	(0.004)	(0.003)
Woman-Year Observations	13.7m	9.3m	13.7m	13.7m	13.7m
Woman FE	Yes	Yes	Yes	Yes	Yes
Province-by-Time FE	Yes	Yes	Yes	Yes	Yes
Province-by-Woman-Age FE	Yes	Yes	Yes	Yes	Yes
Clustering	Prov (29)	Prov (29)	Prov (29)	Prov (29)	Pref (337)

Robust standard errors clustered at the prefecture or province level to allow for autocorrelation of errors over time and space (*** $p<0.01$, ** $p<0.05$, * $p<0.1$). Suitability for cash crops is the (standardised) relative suitability of cash crops compared to grains in the county the mother lives. The other variable relate to past fertility history. The sample consists of a panel of married women born between 1945 and 1959 and their fertility events for years between 1973 and 1990. Column 2 restricts this sample to women who have already had one child. The fixed effects allow for province specific paths of fertility by age, province specific trends in overall fertility and time invariant woman specific fertility differences

Table 5: Education and Fertility

	Any Birth Dummy			
	(1)	(2)	(3)	(4)
Suitability for Cash Crops × 1982-1990	0.017*** (0.004)	0.018*** (0.004)	0.018*** (0.005)	0.018*** (0.005)
Mother Attended Junior Middle School or Higher		-0.033*** (0.002)		-0.030*** (0.002)
Father Attended Junior Middle School or Higher			-0.016*** (0.001)	-0.008*** (0.001)
Woman-Year Observations	13.7m	13.7m	12.7m	12.7m
Woman FE	*No*	*No*	*No*	*No*
County FE	Yes	Yes	Yes	Yes
Parent's Registration Status FE	Yes	Yes	Yes	Yes
Parent's Ethnicity FE	Yes	Yes	Yes	Yes
Province-by-Time FE	Yes	Yes	Yes	Yes
Province-by-Mother-Age FE	Yes	Yes	Yes	Yes

Robust standard errors clustered at the prefecture or province level to allow for autocorrelation of errors over time and space (*** $p<0.01$, ** $p<0.05$, * $p<0.1$). Suitability for cash crops is the (standardised) relative suitability of cash crops compared to grains in the county the mother lives. The other variable relate to past fertility history. The sample consists of a panel of married women born between 1945 and 1959 and their fertility events for years between 1973 and 1990 (including both mothers and non-mothers). The fixed effects allow for province specific paths of fertility by age, province specific trends in overall fertility and time invariant county, parental ethnicity and registration specific fertility differences (separate sets of dummies for each parent). All specifications include controls for fertility history as in table 4.

Table 6: Effects on Fertility by Ethnicity

	Any Birth Dummy		
	All	Han	Non-Han
	(1)	(2)	(3)
Suitability for Cash	0.033***	0.035***	0.012
Crops × 1982-1990	(0.011)	(0.011)	(0.012)
Mother-Year Observations	13.7m	12.9m	0.8m
Mother FE	Yes	Yes	Yes
Province-by-Time FE	Yes	Yes	Yes
Province-by-Mother-Age FE	Yes	Yes	Yes

Robust standard errors clustered at the province level (29) to allow for autocorrelation of errors over time and space (*** $p<0.01$, ** $p<0.05$, * $p<0.1$). Suitability for cash crops is the (standardised) relative suitability of cash crops compared to grains in the county the mother lives. The sample consists of a panel of married women born between 1945 and 1959 and their fertility events for years between 1973 and 1990. Column 2 and 3 restrict the sample to Han and non-Han women respectively. The fixed effects allow for province specific paths of fertility by age, province specific trends in overall fertility and time invariant woman specific fertility differences. All specifications include controls for fertility history as in table 4.

Table 7: Effects on Sex Ratios (Baseline Results)

	Girl Dummy				
	(1)	(2)	(3)	(4)	(5)
SCC × No	0.024***	0.024***	0.024***	0.022***	0.024***
Brother × Born 1982-1990	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)
Observations	890908	890908	830972	691516	890908
County-by-Birth Year FE	Yes	Yes	Yes	Yes	Yes
County-by-No Brother FE	Yes	Yes	Yes	Yes	Yes
Province-by-Birth-Year-by-No Brother FE	Yes	Yes	Yes	Yes	Yes
Clustering	Prov (29)	Prov (29)	Prov (29)	Prov (29)	Pref (337)
Mother Controls		Yes	Yes		
Father Controls			Yes		

Robust standard errors clustered at the province or prefecture level to allow for autocorrelation of errors over time and space (*** p<0.01, ** p<0.05, * p<0.1). Suitability for cash crops is the (standardised, county average) ratio of the value of output of the best cash crop to the value of output of the best grain crop. No Brothers is a dummy taking a value of 1 indicating that at birth the individual had no male older siblings. Born after 1982-1990 is a dummy indicating that the individual was born after 1981. SCC is my normalised measure of suitability for cash crops. The sample is all individuals born between 1973 and 1990 with one older sibling in the 1% sample of the 1990 Chinese Population Census excluding Hong Kong, Macao and Tibet as well as a small number of individuals which could not be matched to geographic data and a small number of other exclusions discussed in the main text. Column 4 also excludes children from households where the number of children is not precisely equal to the number of births. Mother and Father Controls include sets of dummies for age, education (7 categories) and industry of employment (316 categories).

Table 8: Effects on Sex Ratios by Ethnicity

	Girl Dummy		
	All	Han	Non-Han
	(1)	(2)	(3)
SCC × No	0.024***	0.024***	0.007
Brother × Born 1982-1990	(0.008)	(0.008)	(0.019)
Observations	890908	806811	84097
County-by-Birth Year FE	Yes	Yes	Yes
County-by-No Brother FE	Yes	Yes	Yes
Province-by-Birth-Year-by-No Brother FE	Yes	Yes	Yes

Robust standard errors clustered at the province level to allow for autocorrelation of errors over time and space (*** p<0.01, ** p<0.05, * p<0.1). Suitability for cash crops is the (standardised, county average) ratio of the value of output of the best cash crop to the value of output of the best grain crop. No Brothers is a dummy taking a value of 1 indicating that at birth the individual had no male older siblings. Born after 1982-1990 is a dummy indicating that the individual was born after 1981. SCC is my normalised measure of suitability for cash crops. The sample is all individuals born between 1973 and 1990 with one older sibling in the 1% sample of the 1990 Chinese Population Census excluding Hong Kong, Macao and Tibet as well as a small number of individuals which could not be matched to geographic data and a small number of other exclusions discussed in the main text. Column 1 include all births, column 2 only Han births and Column 3 only non-Han births.

Table 9: Labour Force Participation (1990 Individual Cross Section)

	In Labour Force Dummy		
	(1)	(2)	(3)
Suitability for Cash Crops	0.005	0.007	0.009
	(0.003)	(0.005)	(0.005)
Suitability for Cash Crops × Female Dummy	0.011	0.011	0.012
	(0.009)	(0.009)	(0.010)
Female Dummy	-0.142***	-0.143***	-0.135***
	(0.016)	(0.016)	(0.015)
Constant	0.879***		
	(0.004)		
N	7.9m	7.9m	7.9m
Province FE		Yes	Yes
Additional Controls			Yes

Robust standard errors clustered at the province level to allow for autocorrelation of errors over time and space (*** $p<0.01$, ** $p<0.05$, * $p<0.1$). Suitability for cash crops is the (standardised, county average) ratio of the value of output of the best cash crop to the value of output of the best grain crop. The sample consists of individuals over age 18 in the 1% sample of the 1990 Chinese Population Census excluding Hong Kong, Macao and Tibet as well as a small number of individuals which could not be matched to geographic data and a small number of other exclusions discussed in the main text. Additional controls are a full set of dummies for age, household size, registration status and level of education.

Table 10: Labour Force Participation (Prefectural DiD)

	Female LFP			*M/F LFP*	
	(1)	(2)	(3)	(4)	(5)
SCC × I(1990)	0.010	0.010	-0.010	0.010	-0.003
	(0.006)	(0.008)	(0.013)	(0.008)	(0.010)
Suitability for Cash Crops (SCC)	-0.003				
	(0.012)				
SCC × I(1990) × I(Men)				-0.011	-0.011
				(0.007)	(0.007)
N	458	458	458	916	916
Prefectures	229	229	229	229	229
Prefecture FE		Yes	Yes		
Prefecure-by-Gender FE				Yes	Yes
Time FE	Yes	Yes			
Time-by-Gender FE				Yes	Yes
Province-by-Time			Yes		Yes
Clustering	Prov (27)	Prov (27)	Prov (27)	Prov (27)	Prov (27)

Robust standard errors clustered at the province level to allow for autocorrelation of errors over time and space (*** p<0.01, ** p<0.05, * p<0.1). Suitability for cash crops is the (prefectural average standardised using the county s.d for comparability with results elsewhere) ratio of the value of output of the best cash crop to the value of output of the best grain crop. The outcome variable is prefecture and gender specific labour force participation among individuals aged over 18 for a consistent set of prefectures. Data derived from the 1982 and 1990 census microdata, it's construction is described in Online Appendix A.

Online Appendices

(Not necessarily for publication.)

A Data Appendices

This section provides additional details on data construction not discussed in the paper.

A.1 Construction of Prefectural Level Data from 1982 and 1990 Census microdata.

Both the 1982 and 1990 Population Census microdata include GB codes indicating each households prefecture of residence. The prefecture is the administrative unit between the province and the county. There are around 330 prefectures in China, the average province consists of 11 prefectures and the average prefecture 7-8 counties.

In constructing the set of prefectures I face two issues. First, the prefectural identifiers in the microdata are different to the GiS data provided for the same census and so I must match the microdata to the geographic data. Second the prefectural boundaries—and identifiers—change over time, so I must identify the prefectures with consistent boundaries over time.

To solve the first issue, I manually produced a concordance between the identifiers provided in the microdata and those from the GiS maps of census boundaries provided by University of Michigan China Data Center (2006).[32] For a number of prefectures I was unable to make a match between the census microdata and the GiS boundaries. In the 1982 data these prefectures accounted for around 2% of the population.

To solve the problem of prefectural boundaries changing over time I took the following steps. First, I calculated the centroids for all prefectures in all years. Second, for each year, I count the number of centroids contained within a prefectures polygon for the other year. I then discard all prefecture where this number is not equal to one. I then merge the 1982 and 1990 datasets together and drop all prefectures for which data do not exist for all three years. This eliminates prefectures which were split or merged as well as counties with large border changes, however some minor border may changes may remain. The correlation between suitability for cash crops based on 1982 borders with that based on 1990 borders for the sample of counties used is 0.9999, which does not indicate the presence of many large border changes!

After both steps, I am left with a total of 229 usable prefectures, where I am able to observe aggregate variables over areas with consistent boundaries (out of a total of 319 identified in the 1982 census microdata and 347 in the 1990 census microdata). This procedure results in the complete exclusion of prefectures from Shanghai and Hainan, so regressions based on this data have only 27 provinces. As elsewhere, Tibet, Hong Kong and Macau are also excluded from the data.

[32] These maps appear to be identical to those created by the Harvard China Historical GiS Project.

B Additional Results

This section provides additional details on empirical findings not discussed in the main paper.

B.1 Relationship between marital status, propensity to have begun childrearing, and suitability for cash crops

As discussed in section 4.2, the higher fertility observed in counties with faster income growth could, in principle, have been due to earlier marriage and earlier commencement of childrearing. The results in table 4, indicated that at least 2/3 of the increase in fertility was on the intensive margin, which we would not expect to be as strongly influenced by early marriage. However, in table A1, I explore whether suitability for cash crops is associated with the marriage or childrearing propensity.

Using the cross section of women in the 1990 population census, panel A indicates that suitability for cash crops is uncorrelated with the propensity to be married amongst young women. Panel B indicates that if anything young women are more likely to be childless (although the estimated coefficients are small). Panels C and D, indicate that there is no relationship between the age of married or childless young women and suitability for cash crops. In the cross section at least, there is no relationship between the age at which women are getting married or starting a family, and my source of variation in incomes. Although it is possible the absence of a cross sectional relationship masks a differential post reform change, the fact that most of the effect on fertility is observed on the intensive margin suggests that changes in the timing of fertility decisions is unlikely to be driving the results.

Interestingly, panels E and F indicate that amongst men there is some change in the timing of marriage. Men in living in counties which enjoyed faster post reform growth in incomes have a slightly higher propensity to be married and appear to get married at a younger age. This is consistent with higher incomes making it easier for young men to accumulate sufficient savings to enter marriage and form a household. However, the estimated effects remain quite small—a one standard deviation increase in suitability is associated with around a two month decline in the average age of married men between 18 and 30—and so are unlikely to have large impacts on virility.

C Additional Tables

Table A1: Probability of Marriage and Age Married

	(1)	(2)	(3)	(4)
Panel A. Never Married Dummy (Women)				
Suitability for Cash	0.005	-0.001	-0.005	-0.004
Crops	(0.007)	(0.005)	(0.004)	(0.003)
Constant	0.379***			
	(0.008)			
N	1527112	1527112	1527108	1955501
Panel B. Never Had A Child Dummy (Women)				
Suitability for Cash	0.011*	0.009*	0.004	0.003
Crops	(0.006)	(0.005)	(0.004)	(0.003)
Constant	0.484***			
	(0.007)			
N	1527112	1527112	1527108	1955501
Panel C. Age of Married Women				
Suitability for Cash	0.094*	-0.046	-0.032	-0.005
Crops	(0.053)	(0.060)	(0.043)	(0.105)
Constant	25.396***			
	(0.074)			
N	946430	946430	946424	1372504
Panel D. Age of Women Who Have Never Had A Child				
Suitability for Cash	0.026	-0.006	-0.010	-0.008
Crops	(0.049)	(0.034)	(0.023)	(0.027)
Constant	20.987***			
	(0.073)			
N	742402	742402	742401	750139
Panel E. Never Married Dummy (Men)				
Suitability for Cash	-0.032***	-0.029***	-0.032***	-0.026***
Crops	(0.008)	(0.007)	(0.007)	(0.006)
Constant	0.532***			
	(0.010)			
N	1610910	1610910	1610908	2078186
Panel F. Age Married Men				
Suitability for Cash	-0.115**	-0.201***	-0.174***	-0.187*
Crops	(0.055)	(0.070)	(0.054)	(0.095)
Constant	25.991***			
	(0.070)			
N	761925	761925	761923	1196832
Province FE		Yes	Yes	Yes
Controls			Yes	Yes
Sample Born	1960-1972	1960-1972	1960-1972	1955-1972

Robust standard errors clustered at province (29) level to allow for autocorrelation of errors over time and space (*** $p<0.01$, ** $p<0.05$, * $p<0.1$). Suitability for cash crops is the (standardised) relative suitability of cash crops compared to grains in the county the mother lives. 'Controls' are a full set of dummies for educational attainment, registration status (to capture migrant status), ethnicity and age (not panels C, D, F). The sample consists of individuals in 1% sample of the 1990 population census bron between 1960- or 1955-1972 as indicated.

Table A2: Effects on Fertility (Definition of Pre and Post Results)

	Any Birth Dummy		
	(1)	(2)	(3)
Suitability for Cash Crops × 1982-1990	0.033***		
	(0.010)		
Suitability for Cash Crops × 1979-1990		0.031***	
		(0.011)	
Suitability for Cash Crops × 1985-1990			0.026***
			(0.009)
Mother-Year Observations	13.7m	13.7m	13.7m
Mother FE	Yes	Yes	Yes
Province-by-Time FE	Yes	Yes	Yes
Province-by-Mother-Age FE	Yes	Yes	Yes

Robust standard errors clustered at province (29) level to allow for autocorrelation of errors over time and space (*** $p<0.01$, ** $p<0.05$, * $p<0.1$). Suitability for cash crops is the (standardised) relative suitability of cash crops compared to grains in the county the mother lives. Controls for fertility history are included as in table 4. The sample consists of a panel of married women born between 1945 and 1959 and their fertility events for years between 1973 and 1990. The fixed effects allow for province specific paths of fertility by age, as well as province specific trends in overall fertility.

Table A3: Effects on Sex Ratios (Definition of Pre and Post Results)

	Girl Dummy			
	(1)	(2)	(3)	(4)
SCC × No Brother × Born 1979-1990	0.022***			
	(0.008)			
SCC × No Brother × Born 1982-1990		0.024***		
		(0.008)		
SCC × No Brother × Born 1985-1990			0.028***	0.032***
			(0.008)	(0.010)
Observations	890908	890908	890908	606006
County-by-Birth Year FE	Yes	Yes	Yes	Yes
County-by-No Brother FE	Yes	Yes	Yes	Yes
Province-by-Birth-Year-by-No Brother FE	Yes	Yes	Yes	Yes

Robust standard errors clustered at the province level to allow for autocorrelation of errors over time and space (*** p<0.01, ** p<0.05, * p<0.1). Suitability for cash crops is the (standardised, county average) ratio of the value of output of the best cash crop to the value of output of the best grain crop. No Brothers is a dummy taking a value of 1 indicating that at birth the individual had no male older siblings. Born after 1982-1990 is a dummy indicating that the individual was born after 1981. SCC is my normalised measure of suitability for cash crops. The sample is all individuals born between 1973 and 1990 with one older sibling in the 1% sample of the 1990 Chinese Population Census excluding Hong Kong, Macao and Tibet as well as a small number of individuals which could not be matched to geographic data and a small number of other exclusions discussed in the main text. Column 4 excludes individuals born during the reform years (1979-84).