

Bounding Fertility Elasticities

Anson Zhou*

March, 2023

Abstract

I propose a technique for bounding the magnitude of fertility responses to changes in the cost of children, i.e., fertility elasticity, for any country and year. The range is consistent with empirical estimates and is more precise than current meta-analyses.

JEL classification: J13, D13

Keywords: fertility elasticity, pro-natal policies

*National Bureau of Economic Research. 1050 Massachusetts Ave, Cambridge, MA 02138. Email: anzhou@nber.org. I thank the editor and an anonymous referee for helpful comments and suggestions. All errors are my own.

1 Introduction

Just like other elasticities in an economy, the price elasticity of fertility demand is an important moment for economists and policymakers. For economists working with endogenous fertility models, fertility elasticity provides a fundamental discipline to the parameters. For economists interested in child-related policies such as the Child Tax Credit (CTC) and the Earned Income Tax Credit (EITC), fertility elasticity informs the degree of fertility responses. For governments that attempt to raise fertility through transfers, it is important to understand how cost-effective these measures are.

Despite the importance of fertility elasticity and a large body of empirical literature estimating it, little consensus has been reached on its quantitative magnitude (Stone (2020)). Estimation based on historical policies has proven to be difficult for several reasons. First, many historical policies are not sizable enough (relative to the cost of children) to induce measurable changes in fertility (Bergsvik et al. (2020)). Second, given that pro-natal policies are usually nationwide or income-dependent, finding a control group is not straightforward (Gauthier (2005)). This difficulty is especially acute because most pro-natal policies are adopted to address contemporaneous or future fertility decline, while it is unclear if countries in the control group also face a similar situation (Castles (2003)). Third, family policies often come in bundles of incentives in excess of lowering the cost of children, such as measures encouraging women’s labor force participation. Therefore, estimating the fertility responses to a policy bundle is not the same as estimating the fertility elasticity. Lastly, even if past estimates were precise, a widely acknowledged issue with the design-based methodology is that changes in time and institutions limit the applicability of past estimates in a different context.

In this paper, I propose a simple theoretical approach to *bounding* fertility elasticities and discuss its implications. In Section 2, I show that under mild assumptions, a bound can be derived after knowing (1) the prevailing fertility rate and the prevailing cost of children, (2) the maximum desired fertility by households, and (3) the cost of children such that households would rather prefer not to have a child. In Section 3, I apply this method to the United States in 2010 and find that a transfer of size between \$7,500 and \$36,000 raises the fertility rate by 0.1 children per woman, a range tighter than conclusions from meta-analyses of past estimates. I also demonstrate that the bound is simple to compute for any country and year. Section 4 shows that this bound imposes additional restrictions on parameters in endogenous fertility models and is robust to considering heterogeneous fertility and costs of children.

This paper contributes to two strands of literature. The first literature empirically esti-

mates fertility elasticities using historical policies (e.g., [Milligan \(2005\)](#), [Laroque and Salanié \(2008\)](#), [Cohen et al. \(2013\)](#), and [González \(2013\)](#) among many others). This paper has the same goal in mind but approaches the question from a different perspective. Rather than exploiting the local perturbation of prices, this paper bounds the local elasticity using the global properties of the demand curve. The second literature studies fertility choices in the context of labor and macroeconomics (e.g., [Becker and Lewis \(1973\)](#), [Barro and Becker \(1989\)](#), [Córdoba and Ripoll \(2019\)](#)). Despite being a fundamental moment, fertility elasticity is rarely targeted primarily due to the challenges in measurement. The bound derived in this paper could inform calibration in this class of models.

2 Theory

Consider an economy populated by representative agents. I denote the Marshallian demand of fertility as $n(p; \mathbf{p}^{\text{other}}, y)$ where n is fertility, p is the cost of children, $\mathbf{p}^{\text{other}}$ is the price (vector) of other goods, and y is the household's lifetime income.

Assumption 1 The Marshallian demand of fertility $n(p; \mathbf{p}^{\text{other}}, y)$ is downward sloping, continuously differentiable, and convex in p .

This assumption is satisfied by most models of endogenous fertility. See [Section 4](#) for examples.

Assumption 2 There exists \bar{p} and \bar{n} such that $n(\bar{p}; \mathbf{p}^{\text{other}}, y) = 0$ and $n(0; \mathbf{p}^{\text{other}}, y) = \bar{n}$.

This is a mild and realistic assumption. For the first equation, an example is to let $\bar{p} = y$. The existence of \bar{n} reflects biological constraints of childbearing or satiation in preferences.

Proposition 1 The fertility response to price around any pair of (n^0, p^0) is bounded by

$$\left. \frac{dn}{dp} \right|_{(n^0, p^0)} \in \left(\frac{n^0}{\bar{p} - p^0}, \frac{\bar{n} - n^0}{p^0} \right). \quad (1)$$

Proof [Figure 1](#) gives an illustration of the proof. The Marshallian demand of fertility is given by curve BAC where the coordinates of B and C are $(0, \bar{p})$ and $(\bar{n}, 0)$ correspondingly. Point A denotes the prevailing fertility and cost of children (n^0, p^0) . The slope of the demand curve at A (black) is bounded by the slope of AC (red) and the slope of AB (blue) under the Mean Value Theorem and the assumption that curve BAC is decreasing, continuously differentiable, and convex.




Figure 1. Essence of the Proof

Whereas estimation using historical policies relies on local perturbations to p around (n^0, p^0) , the method proposed here exploits the global properties of the demand curve to bound the local slope. Given the practical difficulties in identifying and quantifying a local shock to p , the bounding approach provides a valuable alternative. Adopting a large-scale pro-natal policy or a randomized control trial could lead to a more precise estimate for a specific country and year. But even in that case, the bound remains complementary to the design-based approach because it provides a prediction of the program's cost-effectiveness *ex ante* and a check of the results *ex post*.

3 Quantification

In this section, I calculate the bound for the United States, compare the results with prior studies, and show that the bound is simple to calculate for other combinations of country and year.

In the 2010 U.S., the prevailing fertility n^0 is 1.9 and the cost of one child p^0 is \$458,400 for a middle-income household (Córdoba and Ripoll (2019) Table 1) in 2011 dollars. I make an additional assumption on \bar{n} and \bar{p} :

Assumption 3 Set $\bar{n} = 8$. Choose \bar{p} such that an average household with one child lives in poverty for the rest of their lives, i.e.,

$$y - \bar{p} = y^{\text{poverty}}. \quad (2)$$

Section 4.1 provides further discussions of this assumption.

Following Córdoba and Ripoll (2019), I set $y = \$2,083,200$. Using the federal poverty guideline in 2010, I set $y^{\text{poverty}} = \$941,500$. This implies $\bar{p} = \$1,141,700$. Applying Proposition 1, I find that to raise fertility by 0.1 children per woman, the change in p needed is between \$7,500 and \$36,000.

The bound is tighter than meta-analyses of past estimates presented in Stone (2020). He harmonizes past estimates and concludes that “an increase in the present value of child benefits equal to 10% of a household’s (annual) income can be expected to produce between 0.5% and 4.1% higher birth rates” (see Figure 2). To make the measures comparable, I convert the bound in this paper into percentages using the median annual household income in 2010 (\$49,400). The bound predicts that an increase in the present value of child benefits equal to 10% of a household’s (annual) income can produce between 0.72% and 3.46% higher birth rates.¹

The bound is simple to compute for a different combination of country and year as long we know the prevailing (n^0, p^0) , y , and y^{poverty} . In particular, p^0 acts as a *sufficient statistic* that captures differences in policies, markets, and social norms that affect the costs of child-raising across time and space, while n^0 is the revealed demand under prevailing prices. For example, in the United Kingdom in 2016, the fertility rate was 1.79 children per woman, and the cost of raising a child from birth to 21 years old is £231,800 (CEBR (2016)). The lifetime income y is approximately £1,400,000. p^{poverty} is chosen to be 60% of y (£840,000) following the definition used by the Department of Work and Pensions of the British government. The result shows that to raise fertility by 0.1 children per woman in the United Kingdom, the change in p required is between £3,700 and £18,300.

4 Discussions

In this section, I provide a discussion of the assumption on (\bar{p}, \bar{n}) , the implications of the results, and the robustness under heterogeneous fertility and costs of children.

4.1 Assumptions on Intercepts

Proposition 1 suggests that the choice of \bar{n} and \bar{p} is important to the tightness of the bound. More specifically, the bound will be tighter when \bar{p} or \bar{n} is lower.

The choice that $\bar{n} = 8$ in Assumption 3 is probably conservative towards the high end as the prevailing ideal number of children is around 2.5 children per woman (Stone (2018))

¹In Figure 2, there are some estimates that lie outside of the bound derived in this paper. This could arise because these policies were implemented in a different time and country setting.




Figure 2. Past estimates

Notes: This figure presents a summary of fertility elasticities estimated using historical policies. I select policy changes that include universal child benefits and baby bonuses from the summary file in [Stone \(2020\)](#). When there are multiple estimates exploiting the same policy change, I take the average across studies. The dots represent point estimates of fertility responses to a transfer with a net present value that is 10% of a household's annual income. The red intervals correspond to 95% confidence intervals. The two horizontal dashed lines represent the bound suggested by [Stone \(2020\)](#). The two horizontal solid lines represent the bound derived in this paper for the United States in 2010.

in the United States. The choice of \bar{p} is more difficult and controversial. One option is to set $\bar{p} = y$ because it implies that having a child would cost the parents' lifetime income and leave them with no consumption. Such a choice is arguably too extreme and leads to an imprecise bound - the maximum price change required to induce a 0.1 higher fertility is \$85,500 instead of \$36,000 in the benchmark case. A more realistic approach that I take is to lean on the assumption that an average household prefers not having children to a state where they have one child but live in poverty for the rest of their lives. An interesting avenue for future research is to elicit information on \bar{n} and \bar{p} through surveys. To gauge how sensitive the bound is to these choices, Table 1 provides alternative bounds under different values of \bar{n} and \bar{p} .

\bar{n}	\bar{p}/y	lower bound	upper bound
6	0.45	\$11,200	\$25,200
10	0.45	\$6,000	\$25,200
6	0.65	\$11,200	\$47,100
10	0.65	\$6,000	\$47,100

Table 1. Bounds under alternative values of \bar{n} and \bar{p}

Notes: This table presents the changes in the cost of children needed to raise fertility by 0.1 under different values of \bar{n} and \bar{p} .

4.2 Implications for Models with Endogenous Fertility

Consider a model of fertility choice with dynastic altruism following the notations of [Barro and Becker \(1989\)](#). Agents solve

$$U_0 = \max_{c_t, n_t, k_{t+1}} \sum_{t=0}^{\infty} \beta^t \left(\prod_{i=0}^{t-1} n_i \right)^{1-\varepsilon} \cdot \frac{c_t^{1-\sigma}}{1-\sigma}$$

subject to

$$w_t + (1 + r_t)k_t = c_t + n_t(\chi_t + k_{t+1}) \quad \forall t$$

and the initial k_0 . It is assumed that $\beta, \varepsilon, \sigma \in (0, 1)$ to ensure that children are goods and $\varepsilon \leq \sigma$ for the second-order condition to hold.

The Marshallian demand of fertility in this economy is

$$n_t = (\beta(1 + r_{t+1}))^{\frac{1}{\varepsilon}} \left[\frac{\chi_{t-1}(1 + r_t) - w_t}{\chi_t(1 + r_{t+1}) - w_{t+1}} \right]^{\frac{\sigma}{\varepsilon}} \quad (3)$$

where the net cost of children at time t is $p_t \equiv \chi_t(1 + r_{t+1}) - w_{t+1}$.² When p_t falls by 1 percent, fertility n_t increases by $\frac{\sigma}{\varepsilon}$ percent. Relating this elasticity to the bound computed in the previous section, the value of $\frac{\sigma}{\varepsilon}$ should lie between 0.67 and 3.23. Combining with the prior assumptions on σ and ε , the full set of restrictions are

$$0 < \varepsilon \leq \sigma < 1 \quad \text{and} \quad \sigma < 3.23 \cdot \varepsilon. \quad (4)$$

The bound imposes *additional restrictions* on the choice of σ and ε . For example, [Manuelli and Seshadri \(2009\)](#) satisfies these restrictions with $\sigma = 0.62$ and $\varepsilon = 0.35$; [Córdoba \(2015\)](#) considers $\sigma = 0.3$ and $\varepsilon = 0.288$; and [Daruich and Kozłowski \(2020\)](#) uses $\sigma = 0.5$ and $\varepsilon = 0.25$. The restrictions could become more binding under different $(n^0, p^0, \bar{n}, \bar{p})$.

²This fertility demand satisfies Assumption 1 but not Assumption 2. It can be interpreted as an approximation of the true underlying fertility demand around (n^0, p^0) .

4.3 Heterogeneous Fertility and Costs of Children

Fertility and costs of children are heterogeneous across socioeconomic groups. The calculations above omit these heterogeneities and use average statistics. Here, we check the robustness of the results when heterogeneities across groups are considered.

Consider a heterogeneous-agent economy where individuals belong to different groups that face heterogeneous costs of children p_i^0 and make different fertility decisions n_i^0 . Each group also has different cutoffs \bar{p}_i and \bar{n}_i . Applying Proposition 1, the group-specific bound is given by

$$\left. \frac{dn}{dp} \right|_{(n_i^0, p_i^0)} \in \left(\frac{n_i^0}{\bar{p}_i - p_i^0}, \frac{\bar{n}_i - n_i^0}{p_i^0} \right). \quad (5)$$

To provide a concrete example, let i denote low, middle, and high-income groups in the U.S. The information on (p_i^0, n_i^0) can be found in [Córdoba and Ripoll \(2019\)](#) and [Darulich and Kozłowski \(2020\)](#). To quantify \bar{p}_i and \bar{n}_i , I modify Assumption 3 in the following way:

Assumption 3 (modified) Set $\bar{n}_i = 8$ and choose $\bar{p}_i = 0.5 \cdot y_i$ for all i .

This assumption is motivated by the result in Section 3 that $\bar{p}/y \approx 0.5$.

I find that to raise the fertility by 0.1 children per woman, the change in p needed is between \$5,600 and \$9,900 for the low-income group, between \$7,500 and \$30,700 for the middle-income group, and between \$12,300 and \$73,100 for the high-income group. The aggregation of different bounds is challenging. But as a back-of-the-envelope calculation, the average upper and lower bounds across groups are \$8,500 and \$37,900 respectively. This result is similar to the bound provided in Section 3, i.e., between \$7,500 and \$36,000.

4.4 Implications for Models Assuming Exogenous Fertility

The bounding method provides an evaluation of the exogenous fertility assumption in structural models that analyze child-related policies (e.g., [Mullins \(2019\)](#) and [Guner et al. \(2020\)](#)). A real-world policy example is the Child Tax Credit expansion in the American Rescue Plan. The net present value of this expansion is around \$30,000 per child for fully eligible families.

The results in Section 3 suggest that a \$10,000 reduction in the cost of children raises fertility by 0.03 to 0.13 from the baseline level of 1.9 children per woman in the United States in 2010. Therefore, if we know the extent to which child-related policies affect the costs of children, we can bound the direct fertility responses. Combined with past evidence on how child-rearing affects household labor supply and child quality, we can gauge the biases that assuming exogenous fertility impose on the outcome variables.

One caveat here is that many policies also have effects on income and prices of other goods such as education. In that case, the overall fertility effect could be more muted than

what the bound suggests. One needs to exercise caution when applying this method to policies such as income transfers or education subsidies.

References

- Robert J Barro and Gary S Becker. Fertility choice in a model of economic growth. *Econometrica*, pages 481–501, 1989.
- Gary S Becker and H Gregg Lewis. On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81(2, Part 2):S279–S288, 1973.
- Janna Bergsvik, Agnes Fauske, and Rannveig K Hart. Effects of policy on fertility: A systematic review of (quasi) experiments. 2020.
- Francis G Castles. The world turned upside down: below replacement fertility, changing preferences and family-friendly public policy in 21 oecd countries. *Journal of European Social Policy*, 13(3):209–227, 2003.
- CEBR. Raising a child more expensive than buying a house. 2016. [Link](#).
- Alma Cohen, Rajeev Dehejia, and Dmitri Romanov. Financial incentives and fertility. *Review of Economics and Statistics*, 95(1):1–20, 2013.
- Juan Córdoba and Marla Ripoll. The elasticity of intergenerational substitution, parental altruism, and fertility choice. *The Review of Economic Studies*, 86(5):1935–1972, 2019.
- Juan Carlos Córdoba. Children, dynastic altruism and the wealth of nations. *Review of Economic Dynamics*, 18(4):774–791, 2015.
- Diego Daruich and Julian Kozlowski. Explaining intergenerational mobility: The role of fertility and family transfers. *Review of Economic Dynamics*, 36:220–245, 2020.
- Anne H Gauthier. Trends in policies for family-friendly societies. *The New Demographic Regime: Population Challenges and Policy Responses*, pages 95–110, 2005.
- Libertad González. The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *American Economic Journal: Economic Policy*, 5(3):160–88, 2013.
- Nezih Guner, Remzi Kaygusuz, and Gustavo Ventura. Child-related transfers, household labour supply, and welfare. *The Review of Economic Studies*, 87(5):2290–2321, 2020.
- Guy Laroque and Bernard Salanié. Does fertility respond to financial incentives? 2008.
- Rodolfo E Manuelli and Ananth Seshadri. Explaining international fertility differences. *The Quarterly Journal of Economics*, 124(2):771–807, 2009.
- Kevin Milligan. Subsidizing the stork: New evidence on tax incentives and fertility. *Review of Economics and Statistics*, 87(3):539–555, 2005.
- Joseph Mullins. Designing cash transfers in the presence of children’s human capital formation. *Job Market Paper*.^[235], 2019.

Lyman Stone. How many kids do women want? 2018. [Link](#).

Lyman Stone. Pro-natal policies work, but they come with a hefty price tag. 2020. [Link](#).