Baby Busts and Baby Booms: The Fertility Response to Shocks in Dynastic Models

Larry Jones* and Alice Schoonbroodt** *

August 2007

Abstract

After the fall in fertility during the demographic transition, many developed countries experienced a baby bust, followed by the baby boom and subsequently a return to low fertility (BBB event). Demographers have linked these large fluctuations in fertility to the series of 'economic shocks' that occurred with similar timing – the Great Depression, World War II (WWII), the economic expansion that followed and then the productivity slow down of the 1970s. This paper formalizes a more general link between fluctuations in output and fertility decisions, in simple versions of stochastic growth models with endogenous fertility. First, we develop initial tools to address the effects of 'temporary' shocks to productivity on fertility choices. Second, we analyze calibrated versions of these models. Qualitatively, results show that under reasonable parameter values fertility is pro-cyclical. Moreover, following a deviation in fertility (due to a productivity shock, for example) sets off an endogenous cycle even if productivity is on trend thereafter. Using the U.S. BBB event as a laboratory, we find that the elasticity of fertility to shocks lays between 1 and 1.7, while the elasticity to fertility deviations is about -0.9. Depending on the nature of child costs, the model captures between 35 and 62 percent of the pre-WWII baby bust due to the great depression, while endogenous fluctuations in conjunction with the productivity boom in the 1950s and 1960s captures between 48 and 93 percent of the post-WWII baby boom.

** University of Minnesota and Federal Reserve Bank of Minneapolis, ** University of Southampton.
The authors thank the National Science Foundation for financial support. We also thank Michele Boldrin, Henry Siu and participants at SITE 2005 Conference on The Household Nexus and the Macroeconomy for helpful comments.

1 Introduction

After the fall in fertility during the demographic transition, almost the entire 20th century appears to constitute, for the United States and for most other developed countries, a period of unusual deviations in fertility around an otherwise declining trend. First, fertility rates saw an acceleration of the decrease during 1920-1940: the baby bust. This bust was followed by a large upswing during 1940-1965: the baby boom. During the period 1965-1985 the total fertility rate returned to a second slump at even lower levels. Demographers often link these large fluctuations in fertility to the series of economic 'shocks' that occurred with similar timing – the Great Depression, WWII and the economic expansion that followed, and finally the productivity slow down of the 1970's. To economists, this line of argument suggests a more general link between fluctuations in output growth and fertility decisions, of which the baby Bust-Boom-Bust event (BBB) is a particularly stark example.

One complementary hypothesis also asserted in the demography literature is that the baby boom was a 'catching up' phenomenon from a period of relatively low fertility during the Great Depression. (See Whelpton (1954), Goldberg, Sharp, and Freedman (1959), and Whelpton, Campbell, and Patterson (1966)). One difficulty with the 'catching up' hypothesis that needs to be addressed, according to both demographers and economists, is that, in the data, this 'catching up' clearly does not take place for a given woman (see **?** and section 3.1 below for documentation). That is, completed fertility was low for both the women immediately preceding and immediately following the baby boom mothers. So, if there is 'catching up', it is in a dynastic (aggregate) sense, not at the individual level. That is, in a dynastic model, it is quite possible that the dynasty purposefully 'catches up,' although this need not be observed for any particular generation of women. This is a distinction relevant in our analysis below.

These two hypotheses can be combined to yield the following relationship: current fertility is a stable function of productivity levels and the current stock of people in the economy. That is, unusual drops in income (relative to trend) cause fertility to fall, other things equal; vice versa, unusual increases in income (relative to trend) cause fertility to rise. If the increase in income follows a period of below trend growth, the boost to fertility may be even larger because the current stock of people is low compared to the long run 'target' level. Roughly speaking, in this view fertility is a function of the current income shock and of the stock of population, with either long run average income or long run 'trend' growth rates determining the target level of population, or of family size. In this paper we formalize this link by combining simple versions of the stochastic growth model and the dynastic model of endogenous fertility à la Barro-Becker.

First, we find that fertility is indeed pro-cyclical in most cases. Although it seems unlikely that fertility decisions are affected by quarter to quarter fluctuations in productivity (as addressed in the business cycle literature), longer fluctuations, such as the Great Depression where productivity was below trend for about 10 years, the post-war period where it was above trend for about the same number of years, and the extended productivity slow-down in the seventies, are likely to affect parents outlook on their children's future well being, and therefore their current fertility decisions. In the present context, temporary shocks should therefore be understood as extended swings around a very long term trend. Second, we find that if there are (at least) two periods of productive life, off balanced growth current fertility depends negatively on last period's fertility.

We then use these models to approximate the elasticity of fertility to productivity shocks and fertility deviations. We find that thefirst lays between 1 and 1.7 while the second is close to -0.9. After calibrating to U.S. averages, we find that, through the model, deviations in productivity capture between 35 and 62 percent of the pre-WWII baby bust, while endogenous fluctuations in response to the earlier downward deviation and in conjunction with the productivity boom in the 1950s and 1960s captures between 48 and 93 percent of the post-WWII baby boom.

Among recent conjectures relating cyclical movements in income and fertility, perhaps the best known one was advanced by Easterlin (2000). He argues that fertility decisions are based on expected lifetime income relative to material aspirations which are formed in childhood. When expected income is high relative to individual aspirations, fertility is high, and vice versa. For the women born in the very early years of the twentieth century who were making fertility decisions in the late 1920s and during the 1930s, expected lifetime income was relatively low due to the Great Depression – hence, the baby bust. Vice versa, since the baby boom mothers grew up in bad times (the 1930s) and therefore had low material aspirations – while they were making fertility decisions during good times (i.e. lifetime income was high relative to expectation), they had many children. Operationally, low income for today's generation implies high fertility for the next generation, especially if its expected lifetime income is particularly high.[1]

It is not hard to see that this version of the relationship between (relative) income and fertility decisions, as well as the previous one linking above (below) average growth and high (low) fertility are, in fact, 'dynamic variations' of the Malthusian hypothesis, both in spirit and in their substantive predictions (see Malthus (1798)). Recall that, in the traditional Malthusian view, the long run population level is determined by a fixed natural ratio between available economic resources and population size. When income per capita - i.e. labor productivity - increases above this natural level, fertility also increases until the long run ratio is reestablished, and vice versa for periods of economic crisis. The prediction that periods of unusually high mortality, during which population is depleted while economic resources remain unchanged, should be followed by years of above average fertility, and vice versa is not far from the one considered here: periods of very harsh economic conditions, in which per capita income decreases below the natural level, are also periods in which fertility decreases. Cipolla (1962), Simon (1977), and Boserup (1981) are some of the best historical renditions of such a 'generalized' Malthusian view, and to them we refer for the many details we must by force omit in our brief historical overview.

While one can see from the above that the theory being considered here dates back to the very origins of modern economic demography, surprisingly little has been done to formally address the link between productivity shocks and fertility in a stochastic model of optimal fertility choice. This paper aims at filling this gap by investigating the theoretical

[1]Formally, one version of this story would be to assume that there is habit formation in consumption, as in much of the recent literature on asset pricing.

and quantitative implications of this link in a stochastic version of the traditional dynastic model of endogenous fertility (see Becker and Barro (1988) and Barro and Becker (1989)). Formally, in these models the size of population in period t, N_t, plays a role much like the capital stock in a standard stochastic growth model. Analogously, fertility plays the role of investment. This analogy is sometimes imperfect since, for example, in the Barro-Becker rendition of the dynastic growth model, N_t also enters the utility function of the planner. Thus, in truth N_t has features that are a mixture of capital and consumption in the stochastic growth model. Also, if children cost time, there is an opportunity cost component through wages – absent in the standard stochastic growth model with capital – that makes investment in children cheap in bad times and expensive in good times.

Given those provisos, recall the simple intuition from the single sector growth model with productivity shocks. There is a fundamental desire to smooth consumption due to the concavity of the utility function. Because of this, in a period when productivity is lower than average (and, as a result, output is correspondingly low), agents lower investment to smooth consumption. When the shock is high, the opposite occurs. Thus, the growth rate of K is high when the shock is high and low when it is low. In the case where the analogy to N_t in the endogenous fertility models holds, this implies that the growth rate of N_t, i.e., the fertility rate, is high when the shock is high, and low when the shock is low. These first order deviations induced by variations in current productivity, can be either damped or magnified by the particular type of production function one adopts.

In this paper, we take the simplest case by considering a stochastic version of a Barro-Becker type model where we abstract from all other inputs besides pure labor (such as physical and human capital or land). However, people live and are productive for more than one period. Thus there are young and old workers. Young workers are fertile. Now, whenever the ratio of young to old workers is low (high) (i.e., last period's fertility was low (high)), the "depreciation" into retirement is large. Again, using the analogy to capital in standard growth models, high depreciation at the end of this period implies high investment for consumption smoothing purposes. In the case, where the analogy to fertility holds, the *number of births* is particularly high (low) if it was particularly low (high) last period. The *fertility rate* is the large (small) number of births divided by the small (large) number of young workers and is therefore particularly large (small). If there is only one period of working life, the depreciation rate into retirement (or death) is always 100 percent and hence fertility is independent on last period's fertility.

Next, we analyze a version of the model in which population does not enter the dynastic planner's utility function. This requires a particular configuration of parameters in the Barro-Becker type preferences. In this case only total consumption by all members of the dynasty enters the utility function. This specification simplifies the model and reduces it to one which is analogous to a stochastic Ak model (see Jones and Manuelli (1990) and Rebelo (1991)) with capital vintages. Under the additional assumption that the shocks to productivity are *i.i.d.*, we give analytical characterizations of the model for particular cases.

In Section 4 we perform quantitative experiments on this version of the model. To do this, we first calibrate the model parameters to various averages in the data. We then use actual magnitudes of productivity shocks from 1900 to 2000 and compare the predictions of the model in terms of deviations in fertility rates to the data. We consider two alternative specifications of the model: (1) all costs of raising children are goods costs; (2) all costs are time costs. Reality probably lies somewhere inbetween these two extreme cases. Indeed, the time cost may have become more relevant over time as women entered the labor force. Since women tend to be the parent taking care of children, the relevant opportunity cost is only procyclical if the latter are actually working whenever they are not busy raising children. In case (1), the elasticity of fertility to the current shock is 1.04. In the case where all costs are in terms of time, we find that the elasticity is 1.73.

Since during the Great Depression productivity was about 12 percent below trend, we get about one-half of the 21 percent downward deviation in fertility (TFR) in the first case and about two-thirds in the second. For the baby boom, about 40 to 80 percent of 25 percent upward deviation during late 1950s and 1960s are accounted for by the 10 percent productivity boom alone, while the response to low productivity in the 1930s—and its ensuing low fertility—accounts for another 20 percent in both cases.

All of these statements are contingent on a particular way of identifying the 'shocks' to productivity in the TFP time series. And the results will certainly be affected by how this issue is treated. We perform these and other sensitivity in the last section.

2 Data

In this section, we lay out the basic facts about the time paths of productivity and fertility in the U.S. over the 20th century. We begin with the facts pertaining to the growth in productivity using a consistent measure for Total Factor Productivity (TFP) and Labor Productivity (LP) from Chari, Kehoe and McGrattan (2006). As most economists know, this period is one of more or less continued growth in productivity with a few interruptions. The most significant of these is the great depression. Figure 1 shows the natural logarithms of TFP and LP over the period from 1901 to 2000.

The facts about productivity over this period can be described as follows:

1. the continual upward trend;
2. the marked decline below trend that took place in the 1930s and early 1940s;
3. the return to trend in the early 1950s;
4. the significant increase above trend that took place in the 1950s and 1960s;
5. the productivity slowdown since the 1970s.

This timing of the movements of productivity around trend fits well with the movements in fertility seen in the data. Figure 2 shows the time path of the Total Fertility Rate (TFR) and Cohort Total Fertility Rates (CTFR) (by birthyear of mother +23 years) over the period from 1850 to 2000. We have two time series for TFR, which calculates how many children a woman would have over her lifetime if current age-specific fertility rates were to prevail in the

Figure 1: Total Factor Productivity and Labor Productivity, 1901-2000 (1929=100)

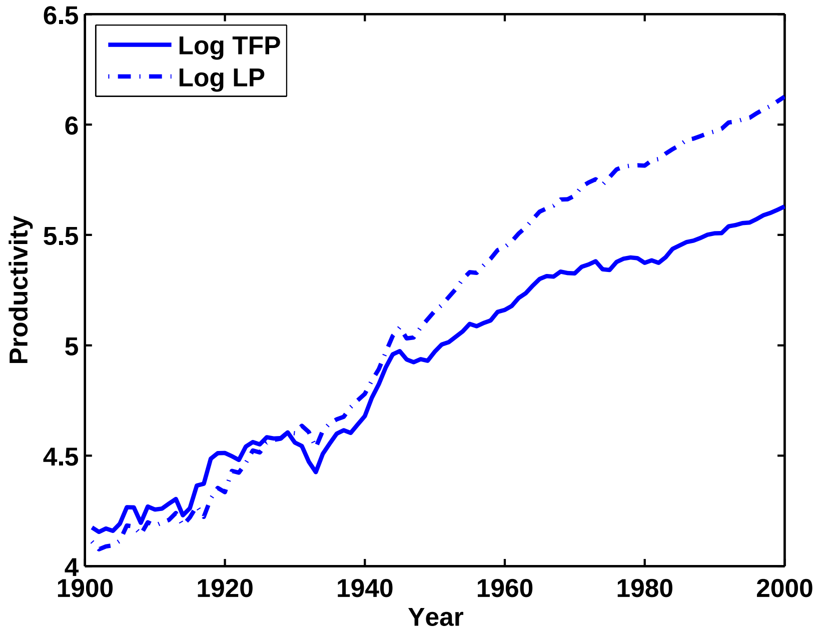

future. The first series is the one prepared by Haines (1994) using Census data and hence is available only every 10 years. The second comes from the Natality Statistics Analysis from National Center for Health Statistics. It is available at annual frequencies, but only since 1917. The CTFR series comes from Jones and Tertilt (2008) and counts how many children were born to a particular cohort of women at the end of their fertile period. Implicitly, it is equivalent to adding age-specific fertility rates pertaining to a particular cohort of women over time. Its frequency is five-year birth cohorts.

At the beginning of the period, fertility is still in the midst of what is known to demographers as the demographic transition, the marked fall in fertility (and mortality) that has occurred in all developed countries. This fall accelerates from the late 1920s to the mid 1930s. Fertility then increases to reach its peak in the baby boom of the 1950s and 1960s. It appears that a good description would be:

1. high, and fairly constantly decreasing fertility from 1850 until 1925, when it reaches a TFR of about 3.5 children per woman;

2. an acceleration of the rate at which fertility is falling between 1925 and 1933 (from TFR=3.5 to TFR=2.1);

3. constant, but low, fertility over the period from 1933 to 1940, with the level at about TFR=2.2;

4. rapidly rising fertility from 1940 to 1957, with TFR going from 2.2 up to 3.7;

Figure 2: (Cohort) Total Fertility Rate, 1850-2000

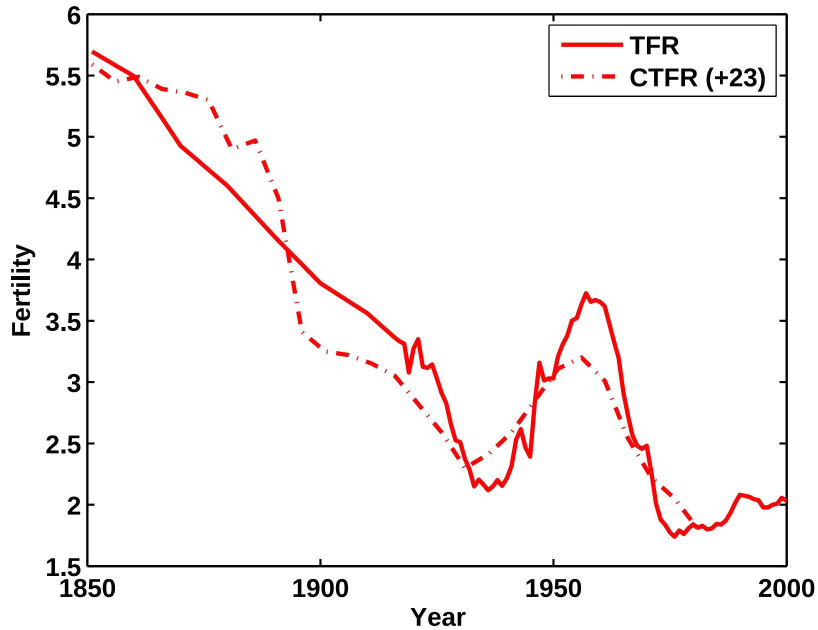

5. high, stable fertility from 1957 to 1961 at about TFR=3.6;
6. a rapid decrease from 1961 to 1976, with TFR going from 3.6 down to 1.7;
7. a slight increase and then stable low fertility over the remainder of the period, with the level at about TFR=2.

We will refer to 2 and 3 as the pre-WWII baby bust, 4 and 5 as the post-WWII baby boom and 6 as the baby bust of the 1970s. The exact sizes of these features of the data depend on how one treats the trend growth in productivity and trend decrease in fertility over the period. For example, was there a common, exogenous growth rate in productivity over the entire period with higher frequency (albeit highly autocorrelated) fluctuations around this trend? Or, were there several regimes of growth? For fertility, one can see that while TFR decreases smoothly over time, the early CTFR data shows that the largest decrease happens for cohorts of women born between 1858 (4.9 children per women) and 1878 (3.25 children per woman). The fluctuations thereafter, however, look very similar in both series, though somewhat larger in TFR than CTFR. This pattern also suggests that if the baby boom is a catching up of fertility from the baby bust, it is at the aggregate (dynasty) level, not at the level of the individual woman.[2]

[2]See also Greenwood, Seshadri, and Vandenbroucke (2005), Doepke, Hazan, and Maoz (2007) and Jones and Tertilt (2008) for more on the make-up, across birth cohorts, of the fertility pattern during the baby boom.

These detrending considerations will have an impact on the analysis we present below and because of this, we try several alternatives. For now, we fit a linear trend to the (ln) TFP and LP series from 1901 to 2000, and detrend TFR using an HP filter. We obtain annual percent deviations over this period plotted in Figure 3.

Figure 3: TFR, TFP and LP Percent Deviations From Trend, 1901-2000

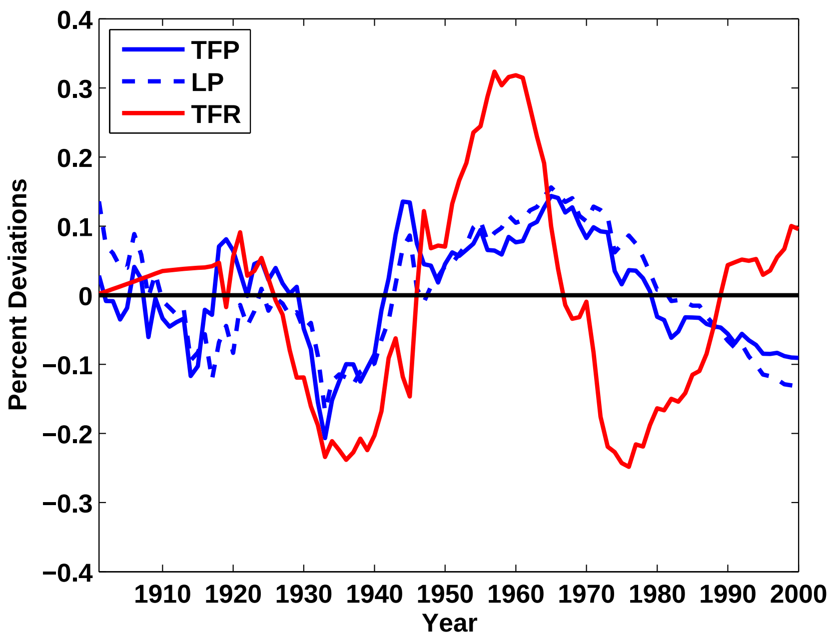

Although it is not perfect, there is an impressive coincidence in timing. The coefficient of correlation between annual TFP and TFR deviations for the years 1901 to 2000 is 0.4, with a coefficient of 0.71 from 1901 to 1940 and a correlation of 0.2 from 1941 to 2000. This suggests that the U.S. TFR is procyclical during the early time period but the correlation is much weaker thereafter (see also Butz and Ward (1979)). As suggested by the model below, one reason for the decrease in the correlation may be the increase in female labor supply which made the opportunity cost of children, women's wages, procyclical.

What our model simulation is also going to capture is the large downward deviation in the 1930s due to the large negative shock during the great depression and a baby boom following endogenously as a response to the baby bust itself, one generation later. We therefore also run the following regression. Let $\widehat{X}_t$ denote the percent deviation from trend in variable X in period t.

$$\widehat{TFR}_t = \lambda_0 + \lambda_1 \widehat{P}_t + \lambda_2 \widehat{P}_{t-l} + \varepsilon_t$$

where $P = \{TFP, LP\}$ and $l = \{20, 25\}$. The results are given in Table 1.

These regressions show that the coefficients on contemporaneous productivity (μ_1) are all positive and significantly different from zero, while the coefficients on productivity a generation ago (λ_2) are all negative and significantly different from zero. Also, λ_1 and λ_2 are of similar magnitude in absolute value, while the constant (λ_1) is likely to be zero.

Table 1: U.S. TFR and Productivity: Regression results

Indep. Var.	Coefficient	Std. error	t-statistic	p-value
Constant	-0.0051	0.0155	-0.3284	0.7435
$\widehat{TFP}_t$	0.8363	0.1851	4.5183	0.0000
$\widehat{TFP}_{t-20}$	-0.8401	0.1951	-4.3068	0.0000
Constant	-0.0112	0.0174	-0.6468	0.5198
$\widehat{TFP}_t$	0.7636	0.2045	3.7329	0.0004
$\widehat{TFP}_{t-25}$	-0.6348	0.2154	-2.9471	0.0043
Constant	-0.0058	0.0169	-0.3395	0.7351
$\widehat{LP}_t$	0.6004	0.1885	3.1845	0.0021
$\widehat{LP}_{t-20}$	-0.6218	0.2009	-3.0952	0.0027
Constant	-0.0139	0.0187	-0.7441	0.4592
$\widehat{LP}_t$	0.4774	0.2252	2.1194	0.0375
$\widehat{LP}_{t-25}$	-0.4796	0.2411	-1.9892	0.0505

2.1 Cross-Country Analysis

TBW

3 The Model

In this section, we lay out a model of the response of fertility to period by period stochastic movements in productivity. To do this, we use a model of fertility based on that developed in Becker and Barro (1988) and Barro and Becker (1989) (Barro-Becker henceforth). The simplification that we make is to assume that there is no physical (or human) capital in the model. Thus, the flow of income is solely due to wage income. On the other hand, we add a stochastic component as well as an explicit age-structure to the basic Barro-Becker model.

3.1 Model setup

A period is 20 years. Every person lives for four periods, one as a child and $T = 3$ as an adult.[3] There is an initial age distribution of the population given by (N_0^3, N_0^2, N_0^1) where N_0^3

[3]The extension to a larger number of (shorter) periods is straightforward, but drastically increases the number of state variables in the model. This complicates the numerical procedures studied below while

is the number of initial old (i.e., their age in period $t = 0$ is $a = 3$), etc. We will normalize by assuming that $N_0^3 = 1$.

No decisions are made in childhood (age $a = 0$). At age $a = 1$, young workers make fertility and consumption decisions and supply one unit of labor inelastically to earn a wage, w_t^1. At age $a = 2$, old workers make consumption decisions and supply one unit of labor inelastically to earn a wage, w_t^2, but are no longer fertile. In the last period of their lives, age $a = 3$, they are retired and only consume.

Adults care about consumption, the number of children and their children's future utility. Following the original Barro-Becker formulation, we assume that the utility of a person who was born in period $t-1$ and whose first period as an adult is in period t is given by:

$$U_t^1 = V_t^1 + \phi\beta g(n_t)U_{t+1}^1$$

where U_t^1 represents the full value of utility of an age $a = 1$ adult in period t looking from that point forward, V_t^1 is the utility this person gets from his own path of consumption, n_t is the number of children that he has and U_{t+1}^1 is the utility that his typical child will receive.

We distinguish between time preference as measured by the discount factor, β and the degree of altruism between generations within a period, ϕ. I.e., $\phi = 1$ means that a person cares as much about the utility of his children as he cares about his own (see, Manuelli and Seshadri (Forthcoming)).

Let c_s^a be the amount of consumption for the typical person in period s that is age a. We will assume that utility from the time path of own consumption $((c_t^1, c_{t+1}^2, c_{t+2}^3)$ – young worker, old worker, retirement) is of the form:

$$V_t^1 = \sum_{a=1}^{3} \beta^{a-1}u(c_{t+a-1}^a).$$

Substituting gives:

$$U_t^1 = \sum_{a=1}^{3} \beta^{a-1}u(c_{t+a-1}^a) + \phi\beta g(n_t)\left[V_{t+1}^1 + \phi\beta g(n_{t+1})U_{t+2}^1\right]$$

Under the assumption that $g(n) = n^\eta$ simplification occurs because $g(n_t)g(n_{t+1}) = g(n_t n_{t+1})$. Thus:

adding little to the intuitions we want to focus on in this paper.

$$
\begin{aligned}
U_t^1 &= \sum_{a=1}^{3} \beta^{a-1} u(c_{t+a-1}^a) + \phi\beta g(n_t) V_{t+1}^1 + (\phi\beta)^2 g(n_t n_{t+1}) U_{t+2}^1 \\
&= \sum_{a=1}^{3} \beta^{a-1} u(c_{t+a-1}^a) + \phi\beta g(n_t) \left[\sum_{a=1}^{3} \beta^{a-1} u(c_{t+1+a-1}^a) \right] + (\phi\beta)^2 g(n_t n_{t+1}) U_{t+2}^1 \\
&= u(c_t^1) + \beta \left[u(c_{t+1}^2) + \phi g(n_t) u(c_{t+1}^1) \right] + \beta^2 \left[u(c_{t+2}^3) + \phi g(n_t) u(c_{t+2}^2) + \phi^2 g(n_t n_{t+1}) u(c_{t+2}^1) \right] + ... \\
&= u(c_t^1) + (\beta\phi) \left[\phi^{-1} u(c_{t+1}^2) + g(n_t) u(c_{t+1}^1) \right] \\
&\qquad + (\beta\phi)^2 \left[\phi^{-2} u(c_{t+2}^3) + \phi^{-1} g(n_t) u(c_{t+2}^2) + g(n_t n_{t+1}) u(c_{t+2}^1) \right] \\
&\qquad + ...
\end{aligned}
$$

Similarly, under the same assumptions, we can define the continuation utility for a person who was a young adult in period t (i.e., he was born in period $t-1$) looking forward from the point when he reaches age a.

Of particular interest is the continuation utility of an adult of age T in period 0 – i.e., the initial old. These are the only agents in the model who care about all agents of all ages in all periods. Utility for these agents, U_0^T, is given by:

$$
U_0^T = \sum_{t=0}^{\infty} (\beta\phi)^t \left[\sum_{a=1}^{T} \phi^{1-a} g(N_t^a) u(c_t^a) \right]
$$

where N_t^a is the number of descendants of the initial age T agent that are of age a in period t. I.e.,

$N_t^1 = n_{t-1} N_{t-1}^1$ is the number of births in period $t-1$;
$N_t^a = N_{t-1}^{a-1}$ for $a = 2, ..., T$;
$N_t^a = 0$ for $a > T$;
$N_0^T = 1$.

To ensure the existence of a balanced growth path, we assume that the cost of children born in period t is in terms of period t consumption but allow this cost to depend on the wage of young workers – $\theta_t(w_t^1)$. Thus, this allows for the two most common ways of modeling child costs: goods costs – $\theta_t(w_t^1) = \theta_t$ – and time costs – $\theta_t(w_t^1) = bw_t^1$. Thus, feasibility in period t is given by:

$$
\sum_{a=1}^{3} N_t^a c_t^a + \theta_t(w_t^1) N_t^b \leq \sum_{a=1}^{2} w_t^a N_t^a \equiv W_t
$$

where $N_t^b = n_t N_t^1 = N_{t+1}^1$ is the total number of births in period t. If the costs are in terms of time, $\theta_t = bw_t^1$, there is an additional constraint, namely that $0 \leq bN_t^b \leq N_t^1$.

We will assume that wages grow at a constant rate, γ, on average. Moreover, there is an aggregate shock so that the entire age specific wage profile shifts up and down by s_t in period t. Thus, in period t, wages are given by:

$$
(w_t^1, w_t^2) = (\gamma^t s_t w^1, \gamma^t s_t w^2).
$$

We assume that s_t is a first order Markov Process (i.i.d. soon).

We will also assume that the costs of children also grow at rate γ. That is, we will assume that, $\theta(w_t^1) = \gamma^t \theta$ in the goods cost case and $\theta(w_t^1) = \gamma^t b w^1$ in the time cost case.

Given this, the Planner's Problem is:

$$P(\gamma, \beta; \{N_0^a\}, s_0) \qquad \max \qquad U_0^T = E\left[\sum_{t=0}^{\infty} (\beta\phi)^t \left[\sum_{a=1}^{T} \phi^{1-a} g(N_t^a(s^{t-a})) u(c_t^a(s^t))\right] |s_0\right]$$

subject to:

$\sum_{a=1}^{3} N_t^a(s^{t-a}) c_t^a(s^t) + \theta(s^t) N_t^b(s^t) \leq \gamma^t s_t \sum_{a=1}^{2} w^a N_t^a(s^{t-a})$;
$N_t^1(s^{t-1}) = N_{t-1}^b(s^{t-1})$ is the number of births in period $t-1$;
$N_t^a(s^{t-a}) = N_{t-1}^{a-1}(s^{t-1-(a-1)})$ for $a = 2, ..., T$;
$N_t^a(s^{t-a}) = 0$ for $a > T$;
N_0^a given, $a = 1, ..., T$.

where $s^t = (s_0, s_1, ..., s_t)$ is the history of shocks up to and including period t.

Let us introduce one additional piece of notation. Let $n_{bt} = \frac{N_t^b}{0.5 N_t^1}$ be the number of births per woman. This is the model quantity that we will identify with the Cohort Total Fertility Rate (CTFR) in the Data, while TFR will be a weighted average over several cohort's fertility.

3.2 Consumption Across the Age Distribution

First, we analyze how consumption within a given period, t, is distributed across the different ages of agents alive at the time. The relevant term in a typical period, t, is:

$$\sum_{a=1}^{3} \phi^{1-a} g(N_t^a) u(c_t^a) = \sum_{a=1}^{3} \phi^{1-a} g(N_t^a) u\left(\frac{N_t^a c_t^a}{N_t^a}\right) = \sum_{a=1}^{3} \phi^{1-a} g(N_t^a) u\left(\frac{C_t^a}{N_t^a}\right),$$

where $C_t^a = N_t^a c_t^a$ is the total consumption of the age a cohort in period t.

Given any level of aggregate consumption in a period, C_t, the planner will choose a distribution across ages to maximize the above subject to $\sum_{a=1}^{3} C_t^a = C_t$. It follows that this is done by equating the marginal utility of a unit of aggregate consumption across the different ages.

To gain more intuition restrict attention to the case where, $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$, and, as above, $g(N) = N^\eta$. For this choice of functional forms, there are two sets of parameter restrictions that satisfy the natural monotonicity and concavity restrictions, both in terms of the aggregate, or dynasty variables, (C, N), and in terms of per capita values, $(N, c) = (N, \frac{C}{N})$. These are: i) $0 < \eta < 1$, $0 < \sigma < 1$ and $0 \leq \eta + \sigma - 1 < 1$, and ii) $\sigma > 1$ and $\eta + \sigma - 1 \leq 0$.[4] The Planner's objective within a period is:

$$\sum_{a=1}^{3} \phi^{1-a} (N_t^a)^{\eta+\sigma-1} \frac{(C_t^a)^{1-\sigma}}{1-\sigma}.$$

As can be seen, the marginal utility of C_t^a is affected by three things – First, is g increasing (case i) or decreasing (case ii) in N ? Second, is N growing over time? Third, what is ϕ?

For example, suppose $\phi = 1$. In this case, the age a term is $g(N_t^a) u \left(\frac{C_t^a}{N_t^a} \right)$. If population is growing (the case typically of empirical interest), then N_t^a is decreasing in a – there are less people in older generations. Thus, if g is increasing (decreasing) in N, $g(N_t^a)$ is decreasing (increasing) in a. Thus, other things equal, the marginal value of an increase in per capita consumption within a period is decreasing (increasing) in age. On the other hand, in this case a given level of aggregate consumption in a cohort is split across fewer people in older groups increasing per capita consumption. This leads to a lower value of u'.

Thus if g is increasing in N, whether the Planner will want consumption to be increasing or decreasing in age within a particular period will depend on which of these two effects is larger. If g is decreasing in N, then if population is growing, per capita consumption is increasing in a for sure. For example, specializing further to the case where $\eta = 1 - \sigma$, and $\phi = 1$, period t utility becomes:

$$\sum_{a=1}^{3} \frac{C_t^{a\,1-\sigma}}{1-\sigma}.$$

Thus, aggregate consumption of all age groups within a period will be equalized, $C^a = C^{a'}$ for all a, a', and hence, larger age groups (younger cohorts if N is increasing) will have smaller per capita consumption – $c_t^a < c_t^{a'}$ for $a < a'$.

3.3 Procyclical fertility and catching up

In this section we study the properties of the solution to the Planner's Problem outlined above. In particular, we will characterize how the policy functions from this problem depend on both the current shock and the initial state.

To gain some intuition about the working of the model, notice that if $\eta + \sigma - 1 = 0$, then N does not enter the period utility function except in aggregate consumption, and hence, if $w^a = 0$, $a \neq 1$, N plays exactly the same role in this model as k does in a stochastic Ak model.[5] There is one twist however. This is that, at least in the case where child-rearing is modeled as a time cost, θ_t – the cost of the investment good – is also stochastic. In that case, since $\theta_t = bw_t^1$, periods when productivity is high are also those when children – the analog of the investment good in the Ak model – are expensive. Also, in this case aggregate consumption, C, grows at the same rate as N (if $\gamma = 1$), but per capita consumption is constant (without shocks). Other than that, the analogy is very close.

[4]The knife-edge case where $\sigma \to 1$ and $\eta = \delta(1 - \sigma)$ can also be analyzed. See Jones and Schoonbroodt (Forthcoming) or **?** for details.

[5]See Jones and Manuelli (1990) and Rebelo (1991), seminal papers on this model.

As is usually true in models with exogenous, trend growth, solutions can be obtained by solving a related model with no growth and a different discount factor. Thus, the solution to $P(\gamma, \beta; \{N_0^a\}, s_0)$ can be obtained directly from the solution to $P(1, \hat{\beta}; \{N_0^a\}, s_0)$ – i.e., $\gamma = 1$ (no growth) and the discount factor, $\hat{\beta}$, depends on γ, β, σ, and η. Because of this result, we will abstract from trend growth through most of the remainder of the paper. In those cases where the solution to the model depends on the discount factor, we will use this result to calibrate to the appropriate discount factor in the detrended model.

Given these standard results, we now derive comparative statics of current fertility with respect to productivity shocks and last period's fertility. To do this, we first simplify the problem to one with only one state variable. We then take first order conditions and analyse comparative statics across steady states/balanced growth paths therein.

Denote by $V(N^1, N^2, N^3; s)$ the maximized value obtained in the problem $P(1, \widehat{\beta}; \{N_0^a\}, s_0)$ when the initial conditions are:

$$(N_0^1, N_0^2, N_0^3) = (N^1, N^2, N^3) \text{ and } s_0 = s.$$

Because of our assumptions on the functional forms for the utility function, it is straightforward to show that the value function is homogeneous of degree η in (N^1, N^2, N^3), i.e., $V(\lambda N^1, \lambda N^2, \lambda N^3; s) = \lambda^\eta V(N^1, N^2, N^3; s)$. The problem $P(1, \widehat{\beta}; \{N_0^a\}, s_0)$ is therefore a standard, stationary dynamic program as long as s is first-order Markov. Because of this result, we can characterize the solution through Bellman's Equation. That is V satisfies:

$$V(N^1, N^2, N^3; s) \equiv \max_{N^{a\prime}, C^a, N^b} \sum_{a=1}^{3} \phi^{1-a} (N^a)^{\eta+\sigma-1} \frac{(C^a)^{1-\sigma}}{1-\sigma} + \beta\phi E\left[V(N^{1\prime}, N^{2\prime}, N^{3\prime}; s')|s\right]$$

$$\text{s.t.} \qquad \sum_{a=1}^{3} C^a + \theta N^b \leq s \sum_{a=1}^{2} w^a N^a;$$

$$N^{a\prime} = N^{a-1}, \text{ for } a = 2, 3;$$
$$N^{1\prime} = N^b.$$

Under the assumptions that $\eta = 1 - \sigma$, and $\phi = 1$, as discussed in Section 3.2, $C^a = C^{a\prime}$ and this problem simplifies to:

$$V(N^1, N^2, N^3; s) \equiv \max_{N^{a\prime}, C, N^b} T \frac{C^{1-\sigma}}{1-\sigma} + \beta\phi E\left[V(N^{1\prime}, N^{2\prime}, N^{3\prime}; s')|s\right]$$

$$\text{s.t.} \qquad TC + \theta N^b \leq s \sum_{a=1}^{2} w^a N^a;$$

$$N^{a\prime} = N^{a-1}, \text{ for } a = 2, 3;$$
$$N^{1\prime} = N^b.$$

As can be seen from this, since $w^3 = 0$ by assumption, we have that $V(N^1, N^2, N^3; s) = V(N^1, N^2, N^{3*}; s)$ for any N^3, N^{3*}. Because of this, we will write V as depending only on $(N^1, N^2; s)$.

Assuming further that shocks are i.i.d., using the homogeneity property of V, the above BE is equivalent to the following one in terms of fertility per household, $n' = N^b/N^1$, with $\widehat{V}(n; s) = V(N^1/N^2, 1; s)/T$, where $n = \frac{N^1}{N^2}$:

$$\widehat{V}(n;s) \equiv \max_{n'} \left[\left(\frac{s\left[w^1 n + w^2\right] - \theta(s) n' n}{T} \right)^{1-\sigma} / (1-\sigma) + \beta n^{1-\sigma} E\left[\widehat{V}(n';s') \right] \right],$$

Taking first order condition with respect to n' in and rearranging gives

$$(FOC) \qquad LHS(n') = \frac{\theta(s)}{TE\hat{V}_1(n',s')} = \left(\frac{s\left[w^1 + \frac{w^2}{n}\right] - \theta(s) n'}{T} \right)^{\sigma} \beta = RHS(n').$$

Then, LHS is increasing in n', with $LHS(0) = 0$ if $E\hat{V}_1(0,s) = \infty$. Also, $RHS(n')$ is decreasing in n' with a positive intercept at $n' = 0$. Thus, there is a unique solution.

To see the behavior of n' as a function of the shock, we consider the two extreme cases, $\theta(s) = \theta$, a goods cost, and a time cost, $\theta(s) = bsw^1$.

In the first case, $\theta(s) = \theta$, the FOC is:

$$(FOC) \text{ Goods cost} \qquad \frac{\theta}{TE\hat{V}_1(n',s')} = \left(\frac{s\left[w^1 + \frac{w^2}{n}\right] - \theta n'}{T} \right)^{\sigma} \beta.$$

In this case, RHS shifts up when s goes up while LHS is unchanged. Thus, n' is increasing in s – fertility is procyclical. This effect is larger, the larger σ. Also, it follows that n' is linear in s.

On the other hand, when $\theta(s) = bsw^1$, a time cost, the FOC becomes

$$(FOC) \text{ Time cost} \qquad \frac{bw^1 s^{1-\sigma}}{TE\hat{V}_1(n',s')} = \left(\frac{w^1(1-bn') + \frac{w^2}{n}}{T} \right)^{\sigma} \beta$$

In this case, if $\sigma > 1$ LHS shifts down when s goes up while RHS is unchanged. Thus, again n' is increasing in s – fertility is procyclical. If, $\sigma < 1$ the opposite occurs and fertility is countercyclical.

Moreover, in both cases, RHS shifts down when the current state, last period's fertility per capita, n, increases while LHS is independent of n. Thus n' decreases when n increases, generating cycles.

Finally, note that if there is only one period of productive life, $w^2 = 0$, FOC becomes:

$$(FOC) \qquad \frac{bAw^1 s^{1-\sigma}}{TE\hat{V}_1(n',s')} = \left(\frac{w^1(1-bn')}{T} \right)^{\sigma} \beta.$$

Hence, in this case, current fertility is independent of last period's fertility.

This leads to the following proposition.

Proposition 1 *Current fertility, $n'(n,s)$ is*

1. a. *procyclical if $\theta(w^1) = \theta$, or if $\theta(w^1) = bsw^1$ and $\sigma > 1$;*
   b. *countercyclical if $\theta(w^1) = bsw^1$ and $\sigma < 1$;*

2. a. *independent of last period's fertility, n if $w^2 = 0$;*
   b. *decreasing in last period's fertility, n if $w^2 > 0$.*

Thus, if $w^2 > 0$ the model generates endogenous cycles, triggered by productivity shocks.

The intuition for why fertility is procyclical in this model is similar to that in many growth models. Here, fertility plays the role of an investment good and the usual consumption smoothing logic implies that when the shock is high, investment should be high so as to offset future negative shocks effects on consumption. This argument is direct when the cost of children is a goods cost. It is tempered when the cost is a time cost by a second effect. This is that the cost of the investment good is also higher than average when the shock is high. Thus, whether or not it is a good idea to invest in those periods depends on how strong the desire is to smooth consumption. When this force is strong – σ is large – the consumption smoothing effect is large relative to the cost effect and fertility is procyclical. Thus, the more important the time cost in raising children, the more procyclical is the cost of children itself. That is, whenever productivity is high, the cost of children is also high and vice versa. This dampens the procyclicality of fertility and, indeed, when the desire to smooth consumption is very low, fertility actually is countercyclical.

Some intuition for the fertility cycles in point 2 of the proposition can also be obtained by analogy with growth models. Here, what we find is a source of endogenous cycles or 'catching-up.' The relevant analogy from capital theory here is to consider a model in which depreciation is not constant. Here, we have an extreme version. Capital (i.e., bodies) that are built in period $t-1$ have full productive capacity in period t and period $t+1$ – there is no depreciation between periods t and $t+1$ – but has zero productive capacity in period $t+2$. Thus, age specific depreciation rates here would be $\delta_1 = 0$ and $\delta_2 = 1$ – no depreciation after one period, full depreciation after two. In a situation like this, when $n = N^1/N^2$ is higher than usual, the planner expects next periods depreciation rate to be lower than usual (because N^2 is relatively low). Because of this, to smooth consumption, current investment – $n' = N^b/N^1$ – will be relatively low. Thus, n' is low when n is high. When there is only one period of working life, the depreciation rate is always 100 percent and hence this effect is not present. Because of this, fertility is independent of last period's fertility. The example with one period working life, $w_t^2 = 0$, is interesting because it corresponds most closely to that of the original Barro-Becker model while the model with more than one period productive life is more realistic.

Thus, according to our theory, while the baby bust in the 1930s may be explained by the Great Depression, the baby boom in the 1950s is first and foremost a response to low fertility in the past: 'catching up.'

4 Quantitative Results

In this section, we use the facts about the time paths of productivity and fertility in the U.S. over the 20th century laid out in Section 2 to perform quantitative experiments on the model. To do this, we calibrate parameters to selected moments of our data. We then use this model to explore two kinds of questions. First, what does the calibrated model say about the size of responses to shocks to productivity – e.g., what is the elasticity of fertility with respect to a productivity shock? In keeping with the theoretical results of the previous section we study both the current response to a shock and also the lagged response one generation later due to the misalignment of the age structure of the workforce. Second, based on the estimated policy functions from the calibrated model we study the predicted response to a productivity shock like those seen in US history – e.g., the Great Depression.

We find that the answers to these quantitative questions also depends on the nature of the costs of children. Because of this, we give results for two alternative specifications. These are: (1) all costs are time costs ($\theta(s) = bw^1 s$); (2) all costs of raising children are goods costs ($\theta(s) = \theta$). As noted above, these two specifications are qualitatively different in that with a time cost, the costs of children are higher in productivity booms than in busts.[6]

We find that the quantitative responses in the model are economically significant in all cases. For example, the elasticity of fertility to a contemporaneous productivity shocks lies between 1 and 1.7,[7] while the elasticity one period later lies between 0.94 and 1.5. Standard recessions will have rather modest effects on fertility, however. Large recessions, such as the Great Depression have important and long lasting effects, however.

We therefore turn to the historical record of the United States and study the predicted response of fertility to the productivity booms and busts over the 20th century. To do this, we must first construct a series of shocks to feed into the model. It is not obvious how to do this realistically. On the one hand, in the model, it is assumed that the shock for the current period is realized at the beginning of the period. Effectively, this means that individuals know, at the beginning of the period, what the sequence of annual shocks will be over the next 20 years. Even with highly correlated shocks at annual frequency, this assumption seems extreme at best. Another alternative is to decrease the length of a period. This necessarily increases the size of the state space. For example, even decreasing the period length to 10 years and maintaining the i.i.d. assumption increases the size of the state space from 2 dimensional to 5.[8] Relaxing the *i.i.d.* assumption which might be required for 10- as opposed to 20-year periods, adds another state variable. In addition, one would have to address timing of births (between age 20-30 versus 30-40) more seriously. This is not an easy problem (see Doepke, Hazan, and Maoz (2007) *** Sommer (2010) *** for examples). Thus, there are technical difficulties with following the strategy of decreasing the period length.

Because of this, we present results for several alternative methods for constructing the relevant series of productivity shocks for TFP and LP. First, since the data shows that over the period 1940 to 1980 at least 60 percent of all births are to women age 20 to 30,[9] we assume that women have all their children between age 21 and 30. As our baseline experiment, we

[6] Reality probably lies somewhere inbetween these two extreme cases. Indeed, the time cost may have become more relevant over time as women entered the labor force. Since women tend to be the parent taking care of children, the relevant opportunity cost is only procyclical if the latter are actually working whenever they are not busy raising children.

[7] As will be seen below, the model response to a shock is smaller when the cost is in terms of time than when it is in terms of goods. Thus, 'between 1 and 1.7' corresponds to 1 with a time cost and 1.7 with a goods cost.

[8] The state space is three-dimensional in our problem: N^1, N^2 and s. Using the homogeneity results, this can be reduced to two state variables, $N = N^1/N^2$ and s. With 10-year periods (and age-groups) the problem becomes six-dimensional: N_b^2, $N^{1,1}$, $N^{1,2}$, $N^{2,1}$, $N^{2,2}$ and s, which using the same methods can be reduced to a five-dimensional problem.

[9] See Vital Statistics of the United States, Table 1-7.

therefore assume that the relevant productivity shock for them is the average one in the data for that 10 year period. We conduct sensitivity analysis on this specification in Section 5.

We also have to choose which productivity measure to use, Total Factor Productivity (TFP) or Labor Productivity (LP). Second, we use TFP in our baseline case and discuss LP in Section 5.

In our baseline experiment, a shock to productivity (or labor income) the size of the Great Depression gives rise to a contemporary baby bust that accounts for 64 to 103 percent of the reduction in CTFR and 39 to 63 percent of the reduction in TFR seen in the data. Further, the model prediction of the lagged response to such a shock is a baby boom. Combined with the productivity boom in the 1950s and 1960s, the predicted size of the baby boom lies between 84 and 150 percent of the actual size of the baby boom observed in the CTFR data and 53 to 92 percent of the one observed in the TFR data.

4.1 Parameterization

Preference Parameters: Throughout, we assume that $\eta = 1 - \sigma$ and set $\sigma = 3$ following Jones and Schoonbroodt (Forthcoming) among others.[10] The discount factor is set to $\beta = 0.96^{20}$ to match an annual interest rate of about 4 percent.

Wage and Productivity Parameters: From the model, wages are given by $w_t^a = s_t \gamma^t w^a$, where w^a is the base wage (in period 0) for workers of age $a = 1, 2$, γ is the trend productivity growth rate for a 20-year period and s_t is the productivity shock which we assume to be *i.i.d.* over time.[11] Further, for computational reasons, it is convenient to assume a functional form for the distribution of productivity shocks, s_t. We assume that $\ln \hat{s}_t \sim N(0, \sigma_s^2)$ where $\hat{s}_t = s_t e^{\sigma_s^2/2}$ so $E(s_t) = 1$. Thus, the parameter values to be determined are w^1, w^2, γ and σ_s.

We normalize $w^1 = 1$ and choose $w^2 = 1.25$. This is in line with life-cycle earnings profiles from Hansen (1993) and Huggett (1996). For example, using these profiles, Andolfatto and Gervais (2008, p.3750) report that the wages of 40 to 60 year old workers are 25 percent higher than those of 20 to 40 year old workers.

The growth rate of productivity, γ, and the standard deviation of productivity shocks, σ_s, are calculated from the TFP series plotted in Figure 1. First, we compute the linear trend in productivity by running the following ordinary least-squares regression:

$$\ln TFP_t = \alpha_0 + \alpha_1 t + \epsilon_t. \tag{1}$$

We find that $\alpha_1 = 0.0159$ and therefore set $\gamma = (e^{\alpha_1})^{20} = 1.0161^{20}$. That is, productivity grows at an average of 1.61 percent per year over the 20th century.[12]

[10] Mateos-Planas (2002), Andolfatto and Gervais (2008) and Scholz and Seshadri (2009) also use a value of $\sigma = 3$. In the Appendix we perform sensitivity analysis with respect to all parameters.

[11] This is a reasonable approximation for long movements in labor productivity across generations, which within a dynasty are 20 years apart.

[12] An alternative detrending method would be to use a Hodrick-Prescott filter, instead. However, it is not clear how to make this method consistent with the model since the growth rate of productivity is not constant. It is also not clear what value to use for the smoothing parameter since we are interested in long fluctuations, rather than quarter-to-quarter or annual deviations from trend. In any case, for low enough

To pin down the value of σ_s several steps are required. First, in our baseline experiment, we use the 10 most fertile years of each cohort within a dynasty, namely age 20 to 30, to determine the productivity shock this cohort's fertility choice is affected by. Thus, the shock we are interested in is $\ln \hat{s}_t = \ln \left(\sum_{t=1}^{10} e^{\epsilon_t} \right) - \mu$ where $\mu = E(\ln(\sum_{t=1}^{10} e^{\epsilon_t}))$. To approximate its standard deviation, σ_s, we assume that ϵ_t follows an AR1 process and estimate

$$\epsilon_t = \rho_0 + \rho_1 \epsilon_{t-1} + \nu_t, \tag{2}$$

where $\nu_t \sim N(0, \sigma_\nu)$, simulate a long series of $\{\epsilon_t\}$, compute a series $\{\ln s_t\}$ and calculate its standard deviation to get $\sigma_s = 0.07$. Of course, the series, $\ln s_t$, so constructed is not exactly normally distributed. However, for all simulations, the Kolmogorov-Smirnov test produces a p-value of about 0.4. This means that, at any reasonable significance level, the test would accept the null hypothesis that the distribution of $\ln s_t$ is normal and hence this should provide a good approximation.

Costs of Children: Given preference and productivity parameters, we then calibrate child costs, $\theta(\cdot)$, to match an annual population growth rate of 0.645 percent per year (see Haines, 1994, Table 1). In the model, population growth corresponds to $\frac{N_b+N^1+N^2}{N^1+N^2+N^3}$. In a steady state with no uncertainty, this is given by the steady state level of fertility choice – the n satisfying $n'(n,1) = n$. Thus, the target is $n'(n,1) = 1.00645^{20}$. This corresponds to a steady state fertility level of 2.27 children per woman. In the goods cost case, we find $\theta = 0.1932$, while with time costs we find $b = 0.1927$. Since the wage for young workers is normalized to 1, these costs imply that it takes about 20 percent of a young worker's income or time to produce a child. This means that that a two person household can at most have ten children.

Summary: These parameters are summarized in Table 2. We describe sensitivity to the parameter values in Section 5.

Table 2: Parameter Values, Baseline

Parameter	σ	β	w^1	w^2	γ	σ_s	θ (goods cost)	b (time cost)
Value	3.00	0.96^{20}	1.00	1.25	1.0161^{20}	0.07	0.1932	0.1927

4.2 Model Impulse Responses

Given the parameter values from the previous section we calculate the decision rules from the model. These can then be used to estimate the model responses to different size productivity

smoothing parameter resulting productivity shocks would typically be smaller. Assuming that the growth rate of productivity is constant in the model but using HP productivity shocks would result in smaller fertility responses in the experiments at the end of this and the next section.

shocks. The results are summarized in Table 3. The rows of the table correspond to the two alternative types of cost structures for children, goods and time. The first two columns report the elasticities (at the steady state) of (a) current and (b) lagged fertility to productivity shocks. The next two columns give the (c) contemporaneous and (d) lagged change in fertility levels that the model predicts will result from a one percent increase in productivity. Column (e) recalls the steady state level of completed fertility (CTFR).

Table 3: Impulse Response (in percent and levels), Baseline

	Percent Deviations		CTFR Levels		
	Initial Period (a)	Lagged Period (b)	Initial Period (c)	Lagged Period (d)	Steady State (e)
Time Cost	1.039	-0.940	2.298	2.253	2.274
Goods Cost	1.741	-1.561	2.314	2.239	2.274

As can be seen from this table, the fertility response to a 1-percent productivity shock generates a 1.74 percent contemporaneous increase in fertility (column (a)) in the goods cost case with a lagged decrease of 1.56 percent one period later (column (b)). In the time cost case, the corresponding percentage changes are a 1.04 percent contemporaneous increase and a 0.94 decrease one period later.

These magnitudes can be compared to the regression results shown in Table 1. In particular, column (a) is to be compared with λ_1 and column (b) with with λ_2. Thus, if the model fit the data perfectly we should see $\lambda_1 \approx 0.83$, and $\lambda_2 \approx -0.84$ (these values are taken from the first row of Table 1). Although the model quantities are not exactly the same, they are remarkably similar.

Comparing the two cases, the initial effect is 70 percent larger in the goods cost case than it is in the time cost case. The reason for this difference is that when child costs are primarily in terms of time there are two offsetting effects. First, when times are good there is a natural tendency to save for the future to equalize marginal utilities. Here, the only way to do this is by increasing family size. On the other hand, since the cost of raising children is positively related to the wage, when times are good, children are relatively more expensive – a force in the opposite direction. When the costs of children are in terms of goods, this second effect is not at work and hence, the overall effect is larger in this case. As can be seen in the table, as the theory predicts, the first effect dominates if the intertemporal elasticity of substitution is low enough ($\sigma > 1$, see Proposition 1).

Further, since in both cases, the ratio of the initial response to the lagged response is about -90 percent, the half-life (in absolute value) of the effect is 7 periods (or generations) in both cases. In particular, the effect 7 periods later is a decrease of 0.8 percent in the goods cost case and 0.5 percent in the time cost case.

To get a sense about the size of these effects note that a one-percent decrease in productivity for a 10-year period roughly corresponds to a recession where GDP is 5 percent below trend for two years. This would decrease fertility by about 0.02 to 0.04 children per woman with a subsequent baby boom of similar size. Thus, the quantitative effects on fertility of a normal sized recession would be quite modest. As we shall see below, the model predicts that a recession the size of the Great Depression gives a much larger response.

4.3 A Historical Episode: The Great Depression

In this section we use the historical record of the actual series of productivity shocks in the U.S. to study the predicted response of fertility to the productivity busts and booms that have occurred over the last 100 years.

We first focus on the dynasty most affected by the great depression, i.e., the cohort making fertility decisions in the 1930s (followed by the 1950s cohort, the 1970s cohort and the 1990s cohort).

Second, we consider other dynasties also affected by the Great Depression but less so. That is, those making fertility choices either in the late 1920s and early 1930s (i.e., between 1926 and 1935, etc.) or in the late 1930s and early 1940s (between 1936 and 1944, etc.). Finally, we consider the dynasty making fertility choices in the 1920s (between 1921 and 1930, etc). By construction, the Great Depression has no effect on the fertility choices of this last dynasty. Thus, we end up with four separate dynasties. Cohorts within these dynasties make overlapping fertility decisions.

Finally, the measure of fertility constructed in this way – using different shocks for different cohorts – corresponds to CTFR since it represents completed fertility for each cohort. We also construct an analog for TFR, i.e., period fertility, by averaging over the two cohorts making their fertility decisions in any five-year period. In doing this, we solve the model separately for the different dynasties and then aggregate.

4.3.1 CTFR: Most Affected Dynasty (M.A.D.)

First, to focus on the effects of the Great Depression, we start with the dynasty making fertility decisions in the 1910s, 1930s, 1950s, 1970s and 1990s. This dynasty is the most affected by the Great Depression. As a benchmark, we assume that the initial state of the dynasty is at its steady state age-distribution. In other words, the dynasty has been facing average shocks ($s = 1$) for a long time so that its current age-distribution is stationary. Define $n^{ss} \equiv n(s) = n(1) = n'(n(1), 1) = n'(n(s), s)$.

For this dynasty, the relevant series of productivity shocks, s_t, is shown in Table 4. The results are shown in Figure 4. With a goods cost, the model predicts that fertility is 19.1 percent below trend during the 1930's; with a time cost the prediction is 11.8 percent below trend. In the data, CTFR actually fell by 18.5 percent. Thus, the model captures a significant fraction of this movement in either case. This is primarily due to the large, negative shock of 12.4 percent during the Great Depression of the 1930s – the dynasty's young-to-old worker ratio is initially almost at its steady state value and hence effects from this are minimal.

Table 4: Productivity shocks (in percent), Baseline (M.A.D.)

Decade/Cohort	1910s	1930s	1950s	1970s	1990s
Productivity Shock	-1.8	-12.43	6.89	3.90	-8.97

Figure 4: Percent Deviations in CTFR, Baseline (M.A.D.)

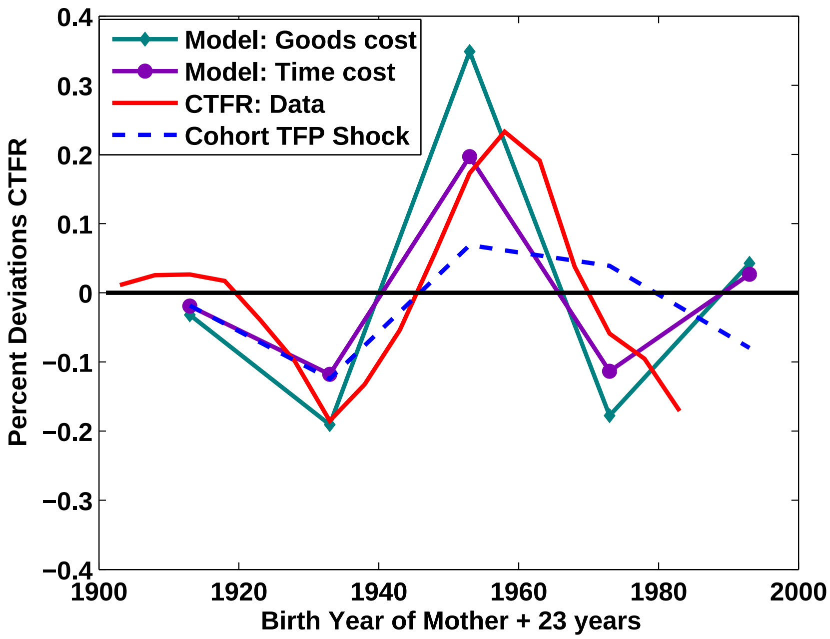

As noted above, there is also a lagged effect in the model from the Great Depression due to the misalignment of the age structure that results from low fertility during the 1930s. This implies, a baby boom for the next cohort in this dynasty which takes place in the 1950s. With a goods cost, the model prediction is an increase in fertility above trend of 34.9 percent, while with a time cost, the corresponding model prediction is 19.7 percent above trend. These model predictions result from a combination of the fertility response to the low young-to-old worker ratio due to the baby bust and the 8.4 percent productivity boom in the 1950s. For the purposes of comparison, the peak of the baby boom in the CTFR data is 23 percent above trend.

In the model, the Great Depression has continuing impacts on the fertility choices in this dynasty – there are continued fluctuations with a decrease of 11 to 18 percent below trend in the 1970s (due to the lagged response to the baby boom but dampened by the 3.9 percent upward fluctuation in TFP in the 1970s); this is followed by an increase of 2.7 to 4.2 percent above trend in the 1990s – again a response to the low young-to-old worker ratio dampened by the negative productivity shock of 8 percent during the 1990s.

This is an extreme version of the baby bust and boom in the model because it is constructed so that the cohort's entire fertile period falls during the Great Depression. For other cohorts or dynasties, the effects are mitigated because at least part of their fertile period is unaffected by the Great Depression. This is also true for the baby boom: the dynasty considered here is most affected by the low young-to-old worker ratio and has a large positive productivity shock in the post-war period.

4.3.2 CTFR: All Dynasties and Cohorts

Table 5 shows how we construct our four dynasties. It illustrates the state variables and shocks for each cohort. For each dynasty, we assume that the initial value of the state variable is $n = n^{ss}$, the balanced growth level of fertility. The first dynasty (d_1) starts with women born between 1876-1880 (labeled 1880), cohort 1 (c_1). They face a productivity shock, s_1, averaging over the period 1901-1910 when they are roughly age 20 to 30. Their fertility choice is $n_{c_1} = n'(n^{ss}, s_1)$. The second dynasty ($d_2$) starts with women born five years later, between 1881-1885 (labeled 1885), cohort 2 (c_2). Their first productivity shock, s_2, is obtained by averaging over 1906-15. Their fertility choice is $n_{c_2} = n'(n^{ss}, s_2)$. The third dynasty ($d_3$) starts with women born between 1886-1890 (labeled 1890), cohort 3 (c_3). This is the most affected dynasty described above. The fourth dynasty (d_4) starts with women born between 1891-1895 (labeled 1895), cohort 4 (c_4). Their first productivity shock s_4, comes from averaging over the period 1916-1925. Cohort 5, born between 1896-1900 (labeled 1900) are the next generation in d_1. Their initial productivity shock s_5, comes from averaging over the period 1921-30. Thus, their fertility choice is given by $n_{c_5} = n'(n_{c_1}, s_5)$. Proceeding further in this fashion leads to Cohort 9, born between 1916-1920 (labeled 1920), the third generation in dynasty 1 making fertility decisions between 1941 and 1950. The shock that this cohort faces is not affected by the Great Depression at all. We expand this series to Cohort 20, born between 1971-1975 (labeled 1975), the third generation in dynasty 1 making fertility decisions between 1996 and 2005.

The fertility choice in the last column of Table 5 is what we plot in Figure 5. Again, we show the results of the model for goods and time costs cases separately. Also plotted is CTFR data from Jones and Tertilt (2008). As expected, we see that dynasties other that the M.A.D. are less affected by the the Great Depression. Still, there is some effect – there are smaller deviations from trend both down and up, 20 years later.

4.3.3 TFR

Next, we compute an analog of TFR, i.e., period fertility, by averaging over the two cohorts making their fertility decisions in any five-year period. That is, we compute TFR in any five-year period by computing the average of fertility of the two cohorts who are making half their fertility decision in that five-year period, weighted by their size (i.e., their parent's fertility choice). For example, $TFR_{1931-35} = \frac{n_{c_6} n_{c_2} + n_{c_7} n_{c_3}}{n_{c_2} + n_{c_3}}$. Note that, if all dynasties are

Table 5: Dynasties and Cohorts, Baseline

Cohort Label	Birthyear Label	Fertile Period	Productivity Shock		State var., n Dynasty				Fertility Choice		
			Label	Value	d_1	d_2	d_3^*	d_4	Label	Goods	Time
c_1	1880	1901-10	s_1	0.50	n^{ss}	...	...	...	n_{c_1}	-1.57	-0.94
c_2	1885	1906-15	s_2	-2.49	...	n^{ss}		...	n_{c_2}	-6.71	-4.06
c_3	1890	1911-20	s_3	-0.45	...		n^{ss}	...	n_{c_3}	-3.20	-1.93
c_4	1895	1916-25	s_4	4.48	...	...	...	n^{ss}	n_{c_4}	5.25	3.12
c_5	1900	1921-30	s_5	3.02	n_{c_1}	...	...	...	n_{c_5}	4.23	2.52
c_6	1905	1926-35	s_6	-5.74	...	n_{c_2}		...	n_{c_6}	-6.20	-3.79
c_7	1910	1931-40	s_7	-11.19	...		n_{c_3}	...	n_{c_7}	-19.05	-11.79
c_8	1915	1936-45	s_8	-0.29	...	...	...	n_{c_4}	n_{c_8}	-7.41	-4.50
c_9	1920	1941-50	s_9	7.21	n_{c_5}	...	...	...	n_{c_9}	6.04	3.55
c_{10}	1925	1946-55	s_{10}	7.14	...	n_{c_6}		...	$n_{c_{10}}$	16.17	9.50
c_{11}	1930	1951-60	s_{11}	8.40	...		n_{c_7}	...	$n_{c_{11}}$	34.88	19.68
c_{12}	1935	1956-65	s_{12}	10.46	...	...	...	n_{c_8}	$n_{c_{12}}$	23.43	13.56
c_{13}	1940	1961-70	s_{13}	12.71	n_{c_9}	...	...	...	$n_{c_{13}}$	13.61	7.93
c_{14}	1945	1966-75	s_{14}	10.45	...	$n_{c_{10}}$		...	$n_{c_{14}}$	1.68	0.75
c_{15}	1950	1971-80	s_{15}	5.37	...		$n_{c_{11}}$	...	$n_{c_{15}}$	-17.76	-11.34
c_{16}	1955	1976-85	s_{16}	-0.17	...	...	...	$n_{c_{12}}$	$n_{c_{16}}$	-19.82	-12.48
c_{17}	1960	1981-90	s_{17}	-3.14	$n_{c_{13}}$	...	...	...	$n_{c_{17}}$	-18.31	-11.31
c_{18}	1965	1986-95	s_{18}	-4.49	...	$n_{c_{14}}$		...	$n_{c_{18}}$	-11.58	-6.83
c_{19}	1970	1991-00	s_{19}	-6.67	...		$n_{c_{15}}$	...	$n_{c_{19}}$	4.25	2.67
c_{20}	1975	1996-00	s_{20}	-3.15	...	...	...	$n_{c_{16}}$	$n_{c_{20}}$	5.10	2.90

Figure 5: Percent Deviations in CTFR, Baseline (All Cohorts)

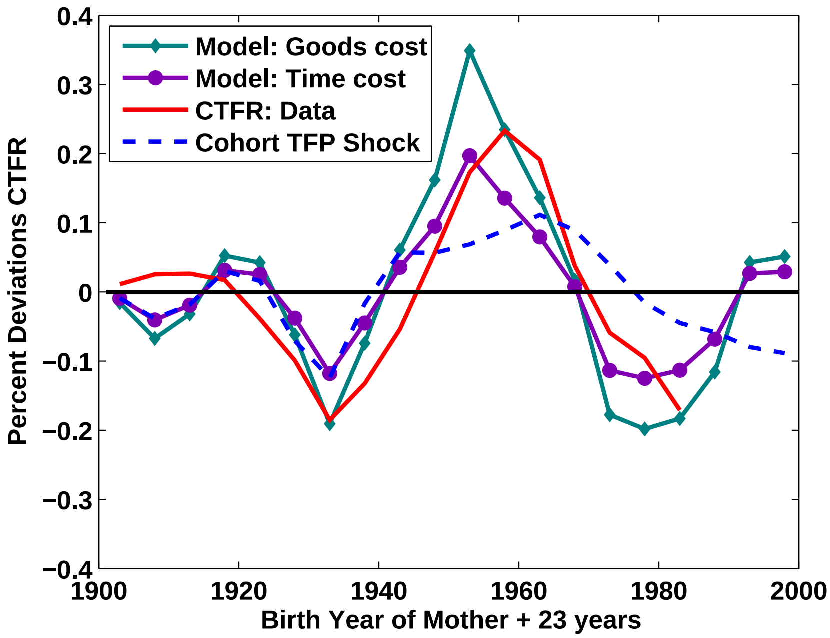

Figure 6: Percent Deviations in TFR, Baseline

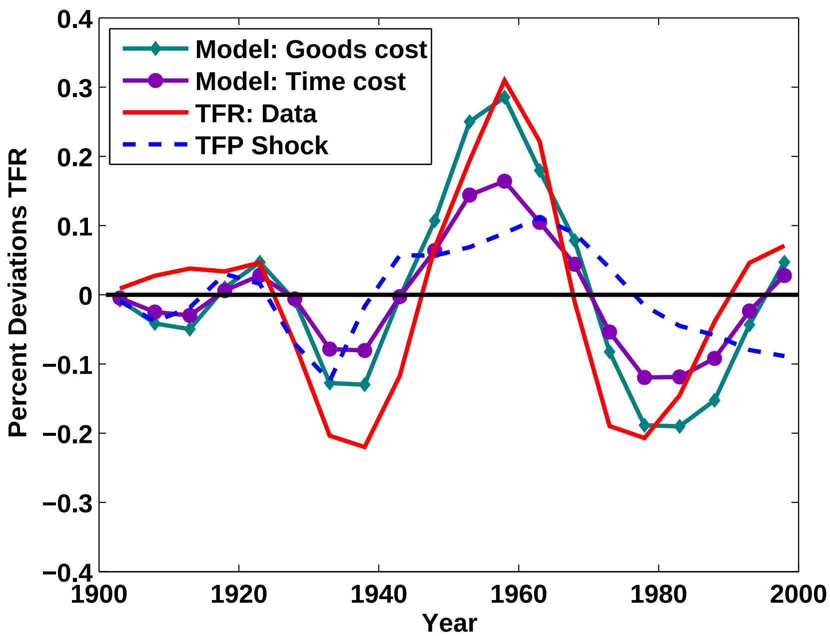

in a steady state $CTFR = TFR = n^{ss}$. One advantage of this fertility measure is that we have data until 2000, while for completed fertility the last cohort in Jones and Tertilt (2008) makes fertility decisions in the 1980s.

The results are shown in Figure 6. Since deviations in TFR in the data are larger than deviations in CTFR and since we are averaging over less affected cohorts, the model accounts for a smaller fraction of these deviations. With a goods cost, the model predicts that fertility would be 12.7 percent below trend in the early and late 1930s; with a time cost, the analog is 7.8 percent. In the data, TFR is 20 percent below trend in the early 1930s and 22 percent below trend in the late 1930s. As with CTFR, this result is mainly due to the negative shock of 11.2 percent during the Great Depression but here the results are mitigated because some women don't make their entire fertility decision during this period. In the data, the peak of the baby boom occurs in the late 1950s when fertility is 31 percent above trend. With a goods cost, the model predicts TFR will be 28.5 percent above trend while with a time cost, the prediction is 16.4 percent. Again, this result is a combination of the fertility response to the low young-to-old worker ratio due to the baby bust and the productivity boom in the 1950s and early 1960s. Hence, the model accounts for 38.5 to 62.6 percent of the baby bust and for 53 to 92 percent of the baby boom in TFR.

In subsequent periods, TFR in the model keeps fluctuating – 11.9 to 18.9 percent below trend in the late 1970s (due to the prior baby boom but dampened by the 3.9 percent increase in TFP). The corresponding movements in the data on TFR is 20.7 percent below trend. The increase in fertility in the 1990s of 7.1 percent above trend in the data is also partly captured by the model which predicts a 2.8 to 4.7 percent increase.

5 Sensitivity

In this section, we discuss the sensitivity of our quantitative results above to (1) changes in the age range used to assign TFP shocks to different cohorts, (2) using labor productivity (LP) instead of TFP and (3) changes in base parameter values.

The first two parts take the parameters in Table 2 as given but change the way productivity shocks are computed. Hence, elasticities and impulse responses remain the same but the size of shocks differs. Results are reported in Figures 8 to 12 and Tables 7 to 9. Increasing the age range tends to decrease the effect of the great depression on fertility because it is diluted across more cohorts. Therefore, the predicted size of the baby boom is also smaller. It also affects the timing of the baby bust and boom. The results using labor productivity are very similar to the baseline, except for the cohort making fertility decisions in the 1910s who experience a productivity boom and those making fertility decisions in the 1920s who experience a significant negative shock. Due to endogenous fluctuations, these changes work themselves through the fertility fluctuations later on in the century.

The third part reports how elasticities and fertility levels change when we change parameter values in Table 2. Results are reported in Table 6. For each parameter change, we report two cases. First, we give results in the case where costs of children remain at their baseline value, which leads to varying fertility levels in steady state. Second, we recalibrate to the target population growth by adjusting costs of children accordingly. Results show that within a reasonable range, parameter choices don't change elasticities and impulse responses very much whether we recalibrate or not. In line with Proposition 1, decreasing the relative wage of old workers, decreases the elasticity to past productivity shocks significantly.

5.1 Age range of productivity shocks

One problem with the assumption that the only relevant shock for the dynasty is when their youngest cohort is age 20 to 30 is that it implicitly assumes that the shock lasts for the entire period of 20 years. In this section, we describe model results when the shock is set to the average productivity shock over the period where the cohort is age 20 to 40.

Figure 7: Cohort TFP Shocks, Age 20-40 versus Age 20-30

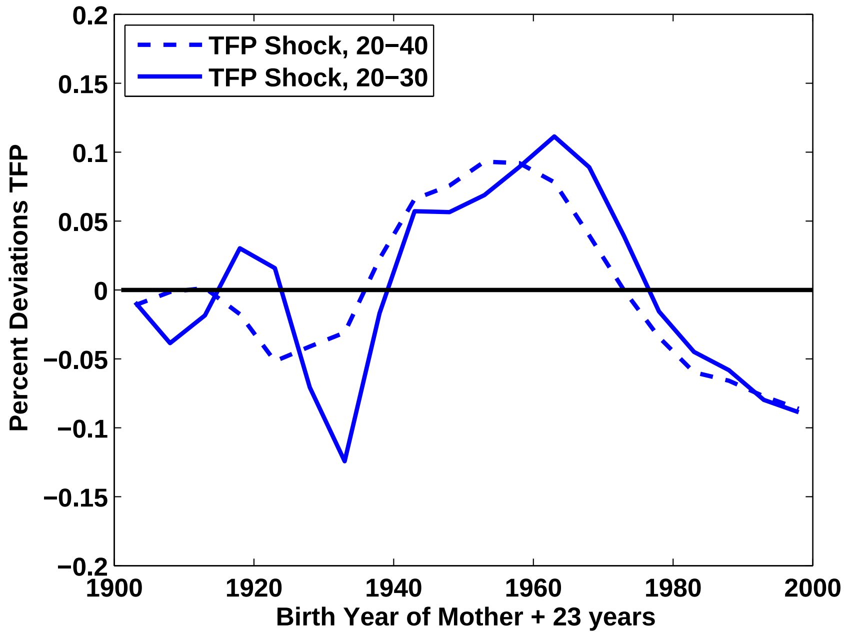

Figure 7 shows cohort productivity shocks for this experiment compared to the baseline. Overall shocks are smoother and smaller because some busts and booms are smoothed out due to the longer fertile period. Also, Cohort c_5 making part of their fertility decisions in the 1920s faces a relatively large negative shock due to the Great Depression in the 1930s when c_5 is age 30 to 40. Therefore, the baby bust starts too early and is significantly smaller than in the baseline model as can be seen in 8 and 9 and in Tables 8 and 9. Because of the smaller baby bust in the 1930s and smaller shocks, the baby boom is also smaller and so are subsequent fluctuations.

The problem with this alternative is that it implicitly assumes that fertility is uniform between age 20 and 40. Hence, the fertility decision of women born at the beginning of the

Figure 8: Percent Deviations in TFR, Age 20-40, Time Cost

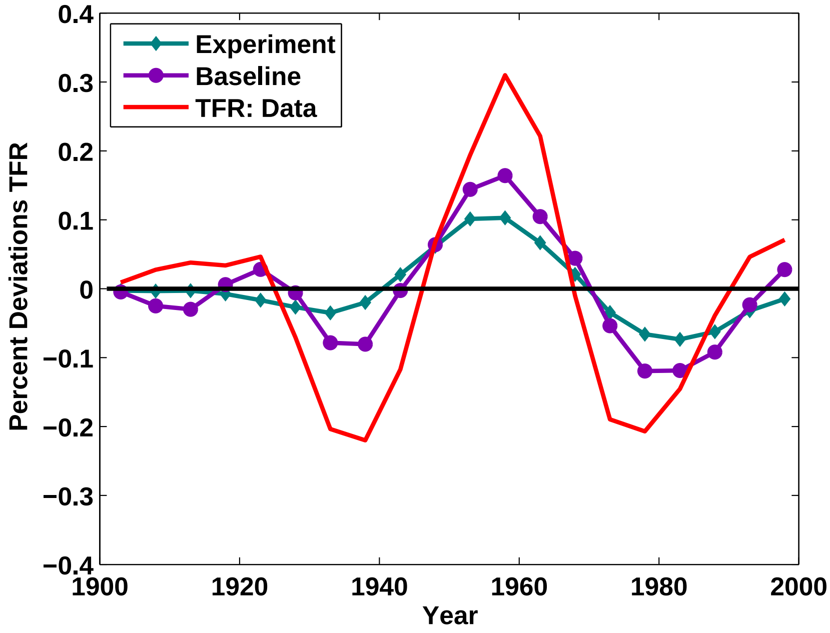

Figure 9: Percent Deviations in TFR, Age 20-40, Goods Cost

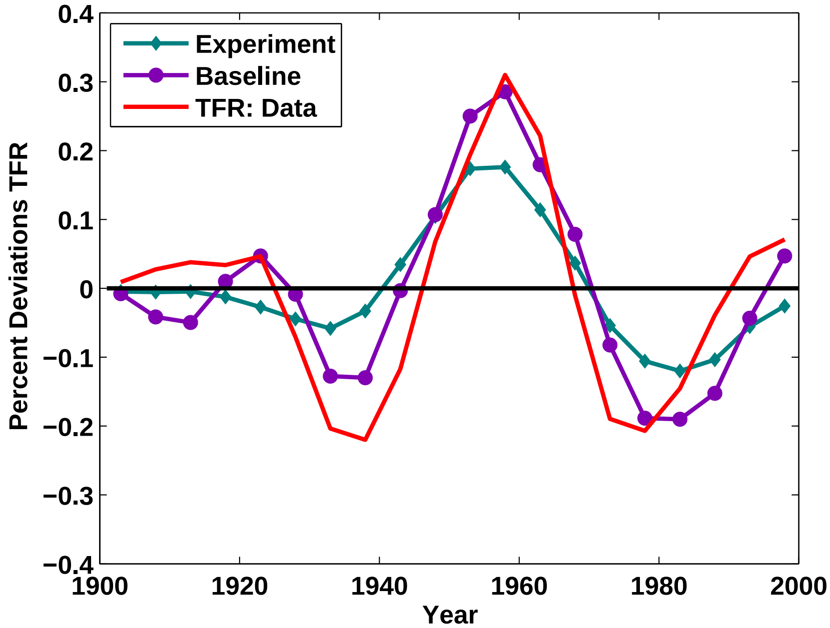

century is heavily affected by the Great Depression in this experiment when in reality most of them had completed their fertility in the 1920s.

Since this alternative is at the opposite extreme of the baseline, the two experiments together should give us reasonable bounds on how much of the fluctuations in fertility over the 20th century in the US the model can account for.

5.2 Labor Productivity versus Total Factor Productivity

Since there is no capital in our model, it is not clear wether we should use labor productivity (LP) or total factor productivity (TFP) to infer productivity shocks. The argument for using TFP is that if capital were included in the model, fertility would respond to TFP alongside investment. The argument for using LP is that most people don't own much capital (besides durable goods and houses) so that their fertility response is to changes in LP.

Figure 10: Cohort TFP Shock, LP versus TFP

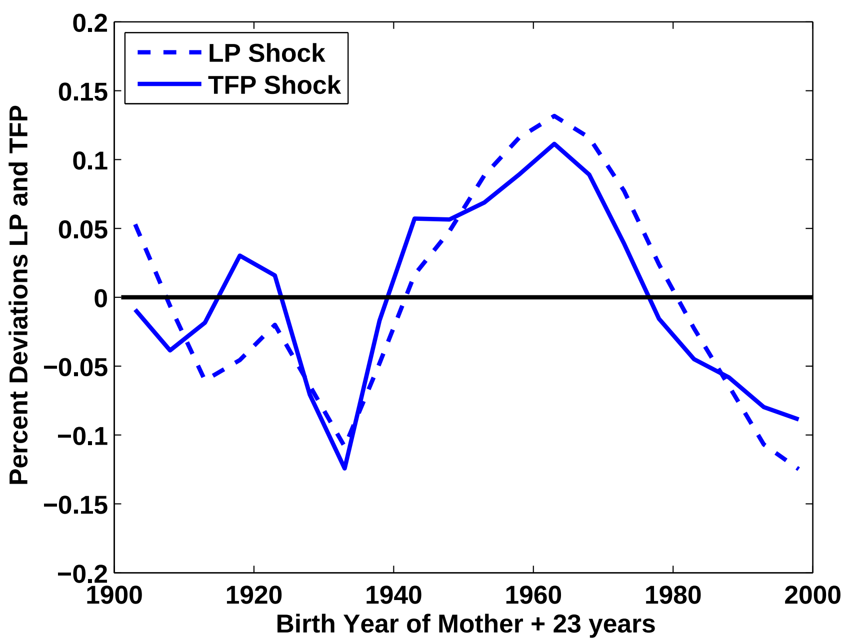

Here we perform the same experiment as in the main text but use deviations from trend in LP instead of TFP. Again, we fit a linear trend through the LP data. The resulting productivity shocks for each cohort are compared to the baseline in Figure 10 and in Table 7. As can be seem in the figure, there are two main differences between the LP compared to the TFP series of shocks. The very first cohort making fertility decisions in the early 1900s experience a productivity boom and those making fertility decisions in the 1910s and 1920s experience a significant negative shock. Therefore, the baby bust is already well underway in the at this time. Despite the similar size of the shock in the 1930s, fertility decreases less in the LP case than in the TFP case. This is because of the mitigating effect of the low young-to-old worker ration due to low fertility in the 1910s. Subsequently, the baby boom starts earlier but is of similar size as the baseline due to the larger LP shock in the 1950s and 1960s compared to TFP. Overall, however, the results are quite similar.

Note that the growth rate of LP in the data is 2.23 percent per year on average, while it is only 1.61 percent per year for TFP. Also, the estimated standard deviation of shocks is 0.08 in most simulations while it is only 0.07 for TFP. In the experiment above, we held these two parameters constant so that we perform only one change at a time. In the next subsection, we perform sensitivity with respect to all parameters, including the productivity growth rate, γ and σ_s, and analyze LP with respect to these changes there.

5.3 Parameters

Next, we compute how elasticities and fertility levels change when we change parameter values in Table 2. Results are reported in Table 6. For each parameter change, we report two cases. First, we give results in the case where costs of children remain at their baseline value, which leads to varying fertility levels in steady state. Second, we recalibrate to the target population growth by adjusting costs of children accordingly. Results show that within a reasonable range of parameter values, the elasticity of contemporaneous fertility to productivity shocks ranges from 1.6 to 1.9 in the goods cost case and 0.95 to 1.2 for the time cost case, while the elasticity of lagged fertility ranges from -1.25 to -1.9 in the goods cost case and -0.7 to -1.2 for the time cost case, whether we recalibrate or not. Thus, the results don't seem to be very sensitive to parameter choices.

A few interesting cases are worth mentioning, however. First, we know from the theory that the choice of σ is important to determine the contemporaneous fertility response to productivity shocks. For a $\sigma = 2.5$, the result is still significant but we know that it tends to zero as $\sigma \to 1$.

Second, since in the detrended version of the model, the effective discount factor is $\beta\gamma^{1-\sigma}$, an increase in β has the same effect as a decrease in γ as long as $\sigma > 1$. As expected, elasticities change accordingly. The more patient people are, the less they respond to productivity shocks. For very low values of γ, of 0.6 percent per year, the recalibration exercise reveals a non-linearity. In this case, the response becomes very large with an elasticity of 2.2 in current fertility and -2.4 one period later. From Jones and Schoonbroodt (Forthcoming) we know that as long as $\sigma > 1$, a decrease in productivity growth generates an increase in fertility. Hence, in the recalibration exercise costs are required to increase significantly (to about 30 percent of a person's time) in order to hit the target fertility level. This may be the reason for the high elasticity.

Third, as mentioned in the previous subsection, LP grows at 2.3 percent per year on average over the 20th century and the estimated standard deviation is 0.8. Table 6 shows elasticities for $\gamma = 1.026^{20}$ and $\sigma_s = 8$. For $\gamma = 1.026^{20}$, the elasticity of current fertility is 1.2 to 1.9 depending on the cost structure and lagged fertility is -1.1 to -1.8. For $\sigma_s = 0.08$, the elasticities are only very slightly higher than those in the baseline. Hence, the effects of LP shocks shown above would be higher if we changed γ and σ_s accordingly.

One important parameter is the relative wage of old to young workers because it de-

Figure 11: Percent Deviations in TFR, Labor Productivity, Time Cost

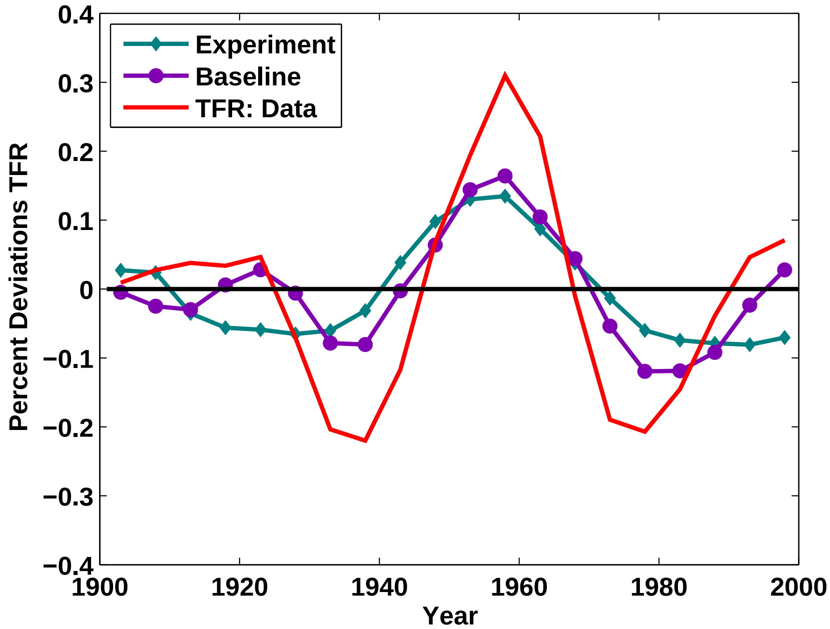

Figure 12: Percent Deviations in TFR, Labor Productivity, Goods Cost

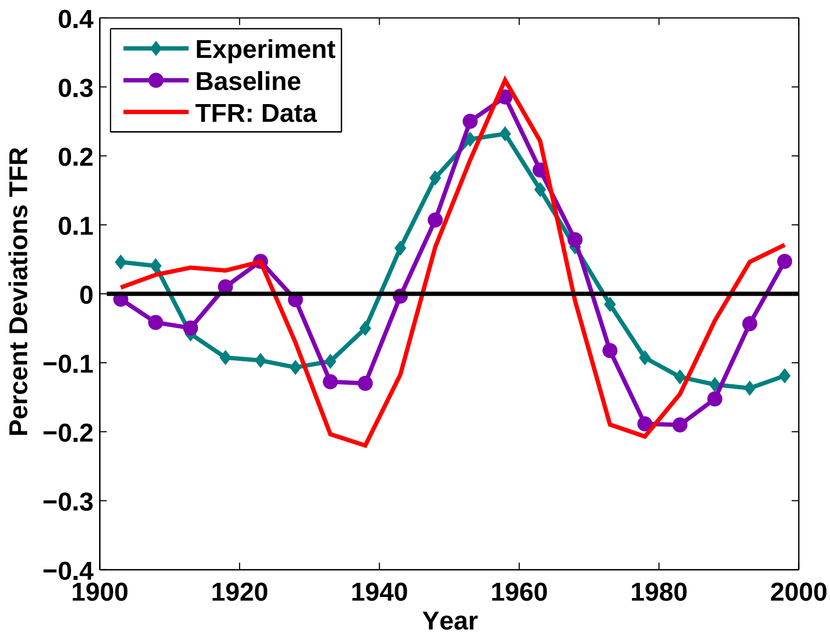

Table 6: Impulse Response (in percent and levels), Parameter Sensitivity

Cost Type	Perc. Dev.		CTFR Lev.			Perc. Dev.		CTFR Lev.		Cost
	Init.	Lag.	Init.	Lag.	SS	Init.	Lag.	Init.	Lag.	(Rec.)
Baseline	($\sigma = 3, \beta = 0.96^{20}, \gamma = 1.016^{20}, \sigma_s = 0.07, w^1 = 1, w^2 = 1.25$)									
Goods	1.74	-1.56	2.31	2.24	2.27	1.74	-1.56	2.31	2.24	0.193
Time	1.04	-0.94	2.30	2.25	2.27	1.04	-0.94	2.30	2.25	0.193
Sensitivity	**No Recalibration**					**Recalibration**				
$\sigma = 3.5$										
Goods	1.81	-1.73	2.17	2.10	2.13	1.83	-1.72	2.32	2.24	0.150
Time	1.16	-1.12	2.16	2.11	2.14	1.20	-1.13	2.30	2.25	0.150
$\sigma = 2.5$										
Goods	1.66	-1.36	2.52	2.44	2.48	1.66	-1.41	2.31	2.24	0.246
Time	0.88	-0.73	2.50	2.46	2.48	0.86	-0.74	2.29	2.26	0.247
$\beta = 0.97^{20}$										
Goods	1.68	-1.40	2.47	2.40	2.43	1.65	-1.40	2.31	2.24	0.242
Time	0.99	-0.83	2.46	2.42	2.44	0.96	-0.82	2.30	2.26	0.243
$\beta = 0.95^{20}$										
Goods	1.81	-1.75	2.15	2.08	2.12	1.82	-1.70	2.32	2.24	0.151
Time	1.09	-1.06	2.14	2.10	2.12	1.11	-1.04	2.30	2.25	0.152
$\gamma = 1.026^{20}$										
Goods	1.88	-1.93	2.03	1.95	1.99	1.88	-1.82	2.32	2.23	0.122
Time	1.14	-1.18	2.02	1.97	2.00	1.17	-1.13	2.30	2.25	0.123
$\gamma = 1.006^{20}$										
Goods	1.61	-1.25	2.63	2.55	2.59	2.22	-2.38	2.32	2.22	0.294
Time	0.94	-0.73	2.62	2.57	2.59	1.29	-1.48	2.30	2.24	0.295
$\sigma_s = 0.08$										
Goods	1.75	-1.57	2.31	2.24	2.27	1.75	-1.57	2.31	2.24	0.193
Time	1.04	-0.94	2.30	2.26	2.28	1.04	-0.94	2.30	2.25	0.194
$\sigma_s = 0.06$										
Goods	1.74	-1.57	2.31	2.23	2.27	1.74	-1.56	2.31	2.24	0.191
Time	1.04	-0.94	2.30	2.25	2.27	1.04	-0.94	2.30	2.25	0.192
$w^2 = 1.5$										
Goods	1.84	-1.88	2.33	2.25	2.29	1.83	-1.87	2.32	2.23	0.198
Time	1.11	-1.15	2.32	2.27	2.30	1.10	-1.14	2.30	2.25	0.199
$w^2 = 1$										
Goods	1.61	-1.21	2.28	2.22	2.25	1.61	-1.20	2.31	2.25	0.186
Time	0.95	-0.71	2.27	2.24	2.25	0.95	-0.71	2.30	2.26	0.187
$w^2 = 0.5$										
Goods	1.29	-0.51	2.23	2.19	2.21	1.30	-0.51	2.30	2.26	0.176
Time	0.73	-0.29	2.23	2.20	2.21	0.73	-0.29	2.29	2.27	0.177

termines the effective depreciation rate into retirement. Since w^1 was normalized to 1, we only move w^2. In the case where $w^2 = 0$, the model exhibits no dependence on last period's fertility because the depreciation rate is the same every period, 100 percent. Table 6 shows how fast the lagged effect disappears as we set w^2 to 1.5, 1 and 0.5.

6 Concluding remarks

TBW

A Appendix

A.1 Transformation of the value function

$$V(N^1, N^2; A) \equiv \max_{N^{1'},C} \qquad T \times C^{1-\sigma}/(1-\sigma) + \beta E\left[V(N^{1'}, N^{2\prime}; s')|s\right]$$

$$\text{s.t.} \qquad T \times C + \theta(s)N^{1'} \leq s\sum_{a=1}^{2} w^a N^a;$$

$$N^{2'} = N^1.$$

The homogeneity property and the fact that $N^{2'}$ is a function of s only implies:

$$\left(N^2\right)^{1-\sigma} V(N^1/N^2, 1; s) \equiv \max_{N^{1'},C} \qquad T\times C^{1-\sigma}/(1-\sigma)+\beta\left(N^{2'}\right)^{1-\sigma} E\left[V(N^{1'}/N^{2\prime}, 1; s')|s\right]$$

$$\text{s.t.} \qquad T \times C + \theta(s)N^{1'} \leq s\sum_{a=1}^{2} w^a N^a;$$

$$N^{2'} = N^1.$$

Letting $c = \frac{C}{N_2}$, we can write this as

$$\left(N^2\right)^{1-\sigma} V(N^1/N^2, 1; s) \equiv \max_{N^{1'},c} \qquad T \times \left(N^2\right)^{1-\sigma}(c^2)^{1-\sigma}/(1-\sigma)$$
$$+\beta\left(N^2\right)^{1-\sigma}\left(\frac{N^{2'}}{N^2}\right)^{1-\sigma} E\left[V(N^{1'}/N^{2'}, 1; s')|s\right]$$

$$\text{s.t.} \qquad T \times N^2 \times c^2 + \theta(s)N^{1'} \leq s\sum_{a=1}^{2} w^a N^a;$$

$$N^{2'} = N^1.$$

Or,

$$\left(N^2\right)^{1-\sigma} V(N^1/N^2, 1; s) \equiv \max_{N^{1'},c} \qquad T \times \left(N^2\right)^{1-\sigma}(c^2)^{1-\sigma}/(1-\sigma)$$
$$+\beta\left(N^2\right)^{1-\sigma}\left(\frac{N^1}{N^2}\right)^{1-\sigma} E\left[V(N^{1'}/N^1, 1; s')|s\right]$$

$$\text{s.t.} \qquad T \times N^2 \times c^2 + \theta(s)\frac{N^{1'}}{N^1}\frac{N^1}{N^2}N^2 \leq s\left[w^1\frac{N^1}{N^2}N^2 + w^2N^2\right].$$

Or,

$$\left(N^2\right)^{1-\sigma} V(N^1/N^2, 1; s)$$
$$\equiv \max_{N^{1'},c} \qquad \left(N^2\right)^{1-\sigma}\left[T \times (c^2)^{1-\sigma}/(1-\sigma) + \beta\left(\frac{N^1}{N^2}\right)^{1-\sigma} E\left[V(N^{1'}/N^1, 1; s')|s\right]\right]$$

$$\text{s.t.} \qquad T \times N^2 \times c^2 + \theta(s)\frac{N^{1'}}{N^1}\frac{N^1}{N^2}N^2 \leq s\left[w^1\frac{N^1}{N^2}N^2 + w^2N^2\right].$$

Clearly, the (N^1, c) that solve the RHS of this identity are unchanged if the $(N^2)^{1-\sigma}$ is removed from the OBJ and the N^2 is removed from the constraint. Thus, the following functional equation must be satisfied by V we have:

$$V(N^1/N^2, 1; s) \equiv \max_{N^{1\prime}, c} \qquad \left[T \times (c^2)^{1-\sigma}/(1-\sigma) + \beta \left(\frac{N^1}{N^2}\right)^{1-\sigma} E\left[V(N^{1\prime}/N^1, 1; s')|s\right]\right]$$

$$\text{s.t.} \qquad T \times c^2 + \theta(s) \frac{N^{1\prime}}{N^1} \frac{N^1}{N^2} \leq s\left[w^1 \frac{N^1}{N^2} + w^2\right].$$

Also, if V solves this then $(N^2)^{1-\sigma} V$ solves the original BE.

Let $n' = \frac{N^b}{N^1} = \frac{N^{1\prime}}{N^{2\prime}}$ and $n = \frac{N^1}{N^2}$ be fertility per household (multiply by 2 to get fertility per woman as in CTFR) as in the original utility function***. Accordingly, define $\widetilde{V}(n; s) \equiv V(n, 1; s)$. Then, it is enough to find the function $\widetilde{V}(n; s)$ that solves:

$$\widetilde{V}(n; s) \equiv \max_{\gamma', c} \qquad \left[T \times (c^2)^{1-\sigma}/(1-\sigma) + \beta n^{1-\sigma} E\left[\widetilde{V}(n'; s')|s\right]\right]$$

$$\text{s.t.} \qquad T \times c^2 + \theta(s) n' n \leq s\left[w^1 n + w^2\right].$$

Assuming that the s are i.i.d. simplifies this in that $E\left[\widetilde{V}(n'; s')|s\right] = E\left[\widetilde{V}(n'; s')\right]$, so that we have:

$$\widetilde{V}(n; s) \equiv \max_{n', c} \qquad \left[T \times (c^2)^{1-\sigma}/(1-\sigma) + \beta (n)^{1-\sigma} E\left[\widetilde{V}(n'; s')\right]\right]$$

$$\text{s.t.} \qquad T \times c^2 + \theta(s) n' n \leq s\left[w^1 n + w^2\right].$$

Let $\widehat{V}(n; s) = \widetilde{V}(n; s)/T$, $\widehat{\theta}(s) = \theta(s)/T$, $\widehat{w}^1 = w^1/T$, $\widehat{w}^2 = w^2/T$ to get:

$$\widehat{V}(n; s) \equiv \max_{n', c} \qquad \left[(c^2)^{1-\sigma}/(1-\sigma) + \beta n^{1-\sigma} E\left[\widehat{V}(n'; s')\right]\right]$$

$$\text{s.t.} \qquad c^2 + \widehat{\theta}(s) n' n \leq s\left[\widehat{w}^1 n + \widehat{w}^2\right].$$

Taking FOC with respect to n' in (be) gives

$$(FOC_{n'}) \qquad (c^2)^{-\sigma} \frac{\theta(s)}{T} n = \beta n^{1-\sigma} E \widehat{V}_1(n', s')$$

$$(BC) \qquad c^2 = \frac{s\left[w^1 n + w^2\right] - \theta(s) n' n}{T}$$

or

$$(FOC_{n'}) \qquad \frac{\theta(s)}{T} n = \left(\frac{s\left[w^1 n + w^2\right] - \theta(s) n' n}{T}\right)^{\sigma} \beta n^{1-\sigma} E \hat{V}_1(n', s')$$

or

$$(FOC_{n'}) \qquad \frac{\theta(s)}{T} n^{\sigma} = \left(\frac{s\left[w^1 n + w^2\right] - \theta(s) n' n}{T} \right)^{\sigma} \beta E \hat{V}_1(n', s')$$

A.2 Sensitivity Tables: Cohort Shocks, CTFR and TFR

Table 7: Productivity Shocks (in percent) for Experiments

	TFP		LP
Cohort	Baseline	Age 20-40	Age 20-30
c_1	-0.90	-1.10	5.30
c_2	-3.85	-0.14	-0.63
c_3	-1.84	0.14	-6.04
c_4	3.02	-1.75	-4.55
c_5	1.58	-5.17	-2.00
c_6	-7.06	-4.11	-6.41
c_7	-12.43	-3.09	-10.84
c_8	-1.68	2.26	-4.79
c_9	5.72	6.59	1.65
c_{10}	5.65	7.57	4.80
c_{11}	6.89	9.31	8.86
c_{12}	8.92	9.21	11.60
c_{13}	11.14	7.81	13.17
c_{14}	8.91	3.95	11.59
c_{15}	3.90	-0.03	7.73
c_{16}	-1.56	-3.43	2.39
c_{17}	-4.49	-5.97	-2.25
c_{18}	-5.82	-6.58	-6.44
c_{19}	-7.97	-7.71	-10.69
c_{20}	-8.86	-8.62	-12.48

Table 8: CTFR response (in percent) from Experiments

	Baseline		TFP, age 20-40		LP, age 20-30	
Cohort	Goods	Time	Goods	Time	Goods	Time
c_1	-1.57	-0.94	-1.92	-1.15	9.22	5.44
c_2	-6.71	-4.06	-0.25	-0.15	-1.09	-0.65
c_3	-3.20	-1.93	0.25	0.15	-10.54	-6.42
c_4	5.25	3.12	-3.05	-1.83	-7.94	-4.81
c_5	4.23	2.52	-7.31	-4.43	-11.05	-6.77
c_6	-6.20	-3.79	-6.94	-4.20	-10.24	-6.23
c_7	-19.05	-11.79	-5.61	-3.39	-9.30	-5.78
c_8	-7.41	-4.50	6.87	4.07	-0.85	-0.56
c_9	6.04	3.55	19.12	11.12	14.36	8.40
c_{10}	16.17	9.50	20.47	11.89	19.23	11.16
c_{11}	34.88	19.68	22.10	12.80	25.56	14.86
c_{12}	23.43	13.56	9.63	5.60	21.02	12.23
c_{13}	13.61	7.93	-2.17	-1.57	9.96	5.66
c_{14}	1.68	0.75	-9.23	-5.85	3.78	1.96
c_{15}	-17.76	-11.34	-16.58	-10.41	-6.55	-4.44
c_{16}	-19.82	-12.48	-13.73	-8.38	-12.08	-7.63
c_{17}	-18.31	-11.31	-8.51	-4.93	-12.02	-7.21
c_{18}	-11.58	-6.83	-2.79	-1.52	-14.36	-8.55
c_{19}	4.25	2.67	3.30	1.94	-12.93	-7.52
c_{20}	5.10	2.90	-1.74	-1.26	-10.79	-6.50

Table 9: TFR response (in percent) from Experiments

	Baseline		TFP, age 20-40		LP, age 20-30	
Time Period	Goods	Time	Goods	Time	Goods	Time
1901-05	-0.78	-0.47	-0.48	-0.29	4.61	2.72
1906-10	-4.14	-2.50	-0.54	-0.32	4.06	2.39
1911-15	-4.96	-2.99	-0.48	-0.29	-5.82	-3.54
1916-20	1.03	0.60	-1.24	-0.75	-9.24	-5.61
1921-25	4.71	2.80	-2.72	-1.65	-9.66	-5.90
1926-30	-0.84	-0.59	-4.43	-2.69	-10.66	-6.51
1931-35	-12.74	-7.84	-5.80	-3.51	-9.79	-6.02
1936-40	-12.99	-8.06	-3.31	-2.01	-5.01	-3.13
1941-45	-0.34	-0.25	3.46	2.06	6.63	3.84
1946-50	10.70	6.38	10.44	6.15	16.81	9.78
1951-55	25.01	14.42	17.37	10.12	22.41	13.02
1956-60	28.54	16.42	17.61	10.28	23.19	13.49
1961-65	17.94	10.47	11.40	6.67	15.10	8.71
1966-70	7.85	4.41	3.67	2.07	6.81	3.77
1971-75	-8.23	-5.37	-5.41	-3.46	-1.52	-1.32
1976-80	-18.86	-11.93	-10.56	-6.60	-9.26	-6.00
1981-85	-19.00	-11.85	-11.99	-7.35	-12.05	-7.43
1986-90	-15.24	-9.19	-10.37	-6.22	-13.15	-7.86
1991-95	-4.35	-2.33	-5.55	-3.20	-13.68	-8.06
1996-00	4.70	2.79	-2.58	-1.49	-11.89	-7.03

References

Andolfatto, D., and M. Gervais (2008): "Endogenous debt constraints in a life-cycle model with an application to social security," *Journal of Economic Dynamics and Control*, 32(12), 3745 – 3759.

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

Becker, G. S., and R. J. Barro (1988): "A Reformulation of the Theory of Fertility," *Quart. J. Econ.*, 103(1), 1–25.

Boserup, E. (1981): *Population and Technological Change*. The University of Chicago Press.

Butz, W. P., and M. P. Ward (1979): "The Emergence of Countercyclical U.S. Fertility," *The American Economic Review*, 69(3), 318–328.

Cipolla, C. M. (1962): *The Economic History of World Population*. Penguin Books.

Doepke, M., M. Hazan, and Y. Maoz (2007): "The Baby Boom and World War II: A Macroeconomic Analysis," NBER Working Paper 13707, National Bureau of Economic Research.

Easterlin, R. A. (2000): "Twentieth-Century American Population Growth," in *The Cambridge Economic History of the United States, Volume 3: The Twentieth Century*, ed. by S. L. Engerman, and R. E. Gallman, pp. 505–48. Cambridge University Press, Cambridge, MA.

Goldberg, D., H. Sharp, and R. Freedman (1959): "The Stability and Reliability of Expected Family Size Data," *Milbank Memorial Fund Quarterly*, 37, 369–85.

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

Haines, M. R. (1994): "The Population of the United States, 1790-1920," Historical Working Paper 56, National Bureau of Economic Research.

Hansen, G. D. (1993): "The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked," *Journal of Applied Econometrics*, 8(1), 71–80.

Huggett, M. (1996): "Wealth distribution in life-cycle economies," *Journal of Monetary Economics*, 38(3), 469–494.

Jones, L. E., and R. E. Manuelli (1990): "A Convex Model of Equilibrium Growth: Theory and Policy Implications," *J. Polit. Economy*, 98(5), 1008–1038.

Jones, L. E., and A. Schoonbroodt (Forthcoming): "Complements versus Substitutes and Trends in Fertility Choice in Dynastic Models," *International Economic Review*.

Jones, L. E., and M. Tertilt (2008): "An Economic History of Fertility in the U.S.: 1826-1960," in *Frontiers of Family Economics*, ed. by P. Rupert, vol. 1. Emerald Press, forthcoming.

Malthus, T. R. (1798): *An Essay on the Principle of Population, as it affects the Future Improvement of Society*. London.

Manuelli, R. E., and A. Seshadri (Forthcoming): "Explaining International Fertility Differences," *Quarterly Journal of Economics*.

Mateos-Planas, X. (2002): "The Demographic Transition in Europe: A Neoclassical Dynastic Approach," *Rev. Econ. Dynam.*, 5(3), 646–680.

Rebelo, S. (1991): "Long-Run Policy Analysis and Long-Run Growth," *J. Polit. Economy*, 99(3), 500–521.

Scholz, J. K., and A. Seshadri (2009): "Children and Household Wealth," Mimeo, University of Wisconsin.

Simon, J. (1977): *The Economics of Population Growth*. Princeton University Press, Princeton, NJ.

Whelpton, P. (1954): *Cohort Fertility*. Princeton University Press, Princeton, NJ.

Whelpton, P., A. Campbell, and J. Patterson (1966): *Fertility Planning in the United States*. Princeton University Press, Princeton, NJ.