

The Elasticity of Intergenerational Substitution, Parental Altruism, and Fertility Choice

Juan Carlos Córdoba*

and

Marla Ripoll†

Revision in progress

Abstract

Dynastic models common in macroeconomics use a single parameter to control the willingness of individuals to substitute consumption both intertemporally, or across periods, and intergenerationally, or across parents and their children. This paper defines the concept of elasticity of intergenerational substitution (EGS), and extends a standard dynastic model in order to disentangle the EGS from the EIS, or elasticity of intertemporal substitution. A calibrated version of the model lends strong support to the notion that the EGS is significantly larger than one, and probably around 2.5. In contrast, estimates of the EIS suggests that it is lower than one. What disciplines the identification is the need to match empirically plausible fertility rates for the U.S.

Keywords: altruism, non-separable utility, intertemporal substitution, intergenerational substitution, imputed value of a child

JEL Classification: D10, D64, D91, J13.

*Department of Economics, Iowa State University, 279 Heady Hall, Ames IA 50010, e-mail: cordoba@iastate.edu.

†Department of Economics, University of Pittsburgh, 4933 W.W. Posvar Hall, Pittsburgh PA 15260, e-mail: ripoll@pitt.edu.

1 Introduction

The infinite horizon model at the core of modern macroeconomics is often motivated as representing in fact a dynasty, a sequence of finitely lived individuals linked by altruism. An implication of this interpretation is that the intra-personal willingness to substitute consumption across periods is the same as the inter-personal willingness to substitute consumption across generations. More precisely, as we formalize below, the model implies that the elasticity of intertemporal substitution (EIS) is identical to the elasticity of intergenerational substitution (EGS). There is, however, no compelling theoretical or empirical reasons why these two parameters, or margins, need to be identical. To the best of our knowledge, we are the first to formally explore the distinction between the EIS and the EGS in macroeconomics.

This paper formally defines the notion of intergenerational substitution and extends an otherwise standard dynastic altruistic model to disentangle the EGS from the EIS. The model illustrates various consequences of isolating the two concepts, and is used to bring empirical discipline to the calibration of the EGS. The analysis yields three main insights. First, while as summarized in Guvenen (2006) estimates of the EIS are commonly below one, we find a much larger EGS of around 2.5 or more. In particular, our calibration features an EIS of 0.67 and an EGS of 2.85. In other words, while the data supports a strong intertemporal consumption smoothing motive, there seems to be much less consumption smoothing across generations.

Second, the shadow price of a child, or the imputed value of a child, plays a key role in identifying the EGS in dynamic altruistic models of fertility choice. We discuss various estimates of the imputed value of a child, mostly calculated as the present value of all the costs of raising a child. This link between the imputed value of a child and the cost of raising him is a direct implication of the equalization of the marginal benefit and the marginal cost of a child from the optimal fertility condition. As we discuss in Section 4.2, according to the USDA (2012) the typical cost of raising a child born in 2011 from ages 0 to 17 for a family of two adults and two children is \$143,051 for a low-income family, \$198,437 for middle-income, and \$328,990 for a high-income family. These figures include direct parental expenses made on children through age 17 such as housing, food, transportation, health care, clothing, child care, and private expenses in education, but abstracts from time costs. Accounting for the time costs of raising children is not trivial, but the best available estimates we can construct suggest that they amount to \$214,576 for a low-income family, \$297,656 for middle-income, and \$493,485 for a high-income family. Therefore, for each of these income groups, the imputed value of a child can be estimated at \$357,627, \$496,093 and \$822,476 respectively. As we show, a large EGS is needed in order to match this range of imputed values of a child. The reason is that the option value of having a child is larger the more inelastic is the willingness to substitute consumption between the parent and the child. We find that if the EGS is lower than one, the inelastic case, then the imputed value of child is much larger, even more than an order of magnitude larger, than what is suggested by the present value cost computation. A similar finding is reported by Murphy and Topel (2006) in a related literature that looks at the value of statistical life for adults. In their case, implausibly large values are obtained when the EIS

is lower than one.

The third insight of our analysis is that the high EGS is also supported by the negative fertility-income relationship documented extensively in the empirical literature.¹ For example, Jones and Tertilt (2008) estimate an income elasticity of fertility of about -0.38 using US Census data. The notable feature of their analysis is that they construct a measure of life-time income by using occupational income and education. Life-time income and fertility are measured for several cross-sections of five-year birth cohorts from 1826-1830 to 1956-1960. They conclude that most of the observed fertility decline in the US can be explained by the negative fertility-income relationship estimated for each cross-section, together with the outward shift of the income distribution over time. The reason why the evidence of a negative fertility-income relationship is supportive of a high EGS is the following. The EGS controls the degree of diminishing returns to lifetime parental income. A low EGS means that parents run into sharp diminishing returns in their own income, and therefore the option value of having a child is larger for richer parents because children provide a way to avoid the decreasing returns. Thus, parents with low EGS will tend to have more children as their income increases. The evidence suggests the opposite and therefore it is supportive of an EGS larger than one.

The dynamic altruistic model of fertility choice we analyze has two key features. First, the utility representation easily allows to associate a single parameter, σ , with the EIS, and a different one, η , with the EGS. Although conceptually very different, the simplicity of our preferences parallels that of Epstein and Zin (1989), and Weil (1990), providing a useful and general framework for analyzing intergenerational issues. While σ is computed from the marginal rate of substitution between consumption in two periods within the lifetime of one individual, η is computed from the marginal rate of substitution between the composite consumption of two generations. In our framework, composite consumption is a CES aggregator of consumption flows within the lifetime of an individual with an elasticity of substitution equal to $1/\sigma$. When $\sigma = \eta$ our framework reduces back to the standard model with additive separability across time and generations.

Second, in our model children are precluded from borrowing and therefore individuals fully depend on parental transfers during childhood. Absent constraints, if a child's future income is larger than the cost of raising him, as the evidence suggests is the case, parents would have incentives to have as many children as possible in order to extract rents from them (see Cordoba and Ripoll, 2014). This would imply not only maximum fertility, but also fertility would be independent of income. Limits on the child's ability to borrow preclude parents from extracting rents, prevent maximum fertility, make possible a negative relationship between income and fertility, and imply strictly positive transfers from parents to children.²

We calibrate the model to income and fertility data across US states. Cross-state data is convenient for our calibration for a number of reasons: a negative relationship between income and

¹See Jones and Tertilt (2008) and Jones, Schoonbroodt and Tertilt (2011) for a recent survey of the literature.

²Less than maximum fertility is possible in the unconstrained version of the model if children are a net financial costs to parents, as in Becker and Barro (1988). However, the evidence reported in Cordoba and Ripoll (2014), Section 3, suggests that the present value of income net of child costs is positive. In other words, children are a net financial benefit to the parents.

fertility is observed across states; the assumption that the interest rate is identical across states can be justified; and a simple representative agent model can still be used. We estimate an income elasticity of fertility of -0.143 , significant at the 5% level on a population-weighted regression. This elasticity is close to the -0.17 estimated by Jones and Tertilt (2008) using individual-level Census data for the most recent cohorts. The key part of the calibration is to select the appropriate targets to provide the identification of three key parameters that capture different aspects of how the utility of the child enters into the utility of the parent. One of these is η , which governs the EGS. Two other parameters enter into the weight that the parent gives to their children. One parameter, α , determines the level of parental altruism while a second parameter, ε , determines the degree of diminishing altruism toward additional children. The model implies that α is the key determinant of the amount of parental transfers. We thus identify α using USDA evidence on the goods costs of raising children. Moreover, ε captures how the altruistic weight changes with the number of children. We select a target that brings us the closest we can to its value by mapping available estimates of the willingness to pay to relieve children’s symptoms from respiratory illnesses, and how this changes with the number of children in the family. Once identification targets are selected for α and ε , we show that η can be identified by requiring the model to match the average US fertility rate. The intuition for the connection between fertility and η is that η is the key determinant of the imputed value of a child and therefore of the incentives to have children. Our benchmark calibration and a number of robustness checks support a value of $EGS = 1/\eta > 1$.

Our analysis suggests that beyond fertility choice, disentangling the EIS from the EGS is important to analyze a number of other issues in modern macroeconomics. Since the EGS captures substitution across generations, the study of longer-term issues such as inequality or any policies that involve intergenerational transfers can be more properly analyzed thinking of the EGS instead of the EIS. In order to illustrate the broader scope of our framework, we provide a simple overview of how disentangling the EIS from the EGS could improve our understanding of US inequality. In Section 5.3 we briefly discuss an intergenerational version of a Bewley model in which the source of inequality is uninsurable idiosyncratic risk in earning ability. Individuals differ not only in their earning ability, but also in the amount of transfers they received from their parents. It is assumed that parents cannot insure their children against their random abilities because transfers are non-contingent, and transfers to children cannot be negative. A preliminary calibration of this model shows that our framework is able to produce more wealth inequality than the standard model in which the EIS and the EGS are not disentangled. The reason is that the standard model, a low EIS (and low EGS) introduces too much aversion to consumption fluctuations, inducing too much savings, which tends to decrease wealth inequality. In contrast, if as we show in this paper the EGS is higher than one, then the prospects of falling into poverty are less painful, or hitting the zero-transfer constraint is less problematic, and therefore a larger fraction of the population ends up hitting the constraint, generating more wealth dispersion. Although mostly a preliminary illustration, this analysis provides a glimpse into the potentially large scope of the framework we propose.

The remainder of the paper is organized as follows. Section 2 uses a simple two-period model to

illustrate the difficulty of standard models to match the imputed value of a child and to motivate the importance of disentangling the EIS from the EGS. The section also introduces the formal definitions of these two distinct concepts. In Section 3 we solve the full version of our dynamic altruistic model of fertility choice. The details of the calibration are presented in Section 4. Section 5 discusses potential extensions of our model, as well as other applications of our framework to modern macroeconomics. Section 6 concludes.

2 The value of a child and the EGS

The purpose of this section is twofold. It first uses a simple two-period model to derive an expression for the shadow price of a child, or the imputed value of a child, and illustrates the limitations of a formulation that does not disentangle the EIS from the EGS. The key issue is that the imputed value of a child is implausibly high for standard values of the EIS below one. The second part of this section extends the basic set up to multiple periods, defines the EGS concept as a separate concept from the EIS, and shows that the imputed value of a child is linked to the EGS rather than the EIS.

2.1 The value of a child in a two-period model

Consider the problem of an altruistic individual who is deciding between having a child or not. Let $U(c) \geq 0$ be the utility associated to consumption flow c . The lifetime utility of an individual with no children and lifetime income y is $V_0 = U(y)$. The lifetime utility of an individual with one child is:

$$V_1 = \max_{c_1 \geq 0, c_2 \geq 0} U(c_1) + \alpha U(c_2) \text{ subject to } c_1 + pc_2 + \lambda y = y + py, \quad (1)$$

where c_1 and c_2 are the consumption of the parent and child respectively, $\alpha \geq 0$ is the degree of altruism, λy is the cost of raising the child where $\lambda \in (0, 1)$, and p is the price of a bond. The budget constraint incorporates the present-value lifetime income of the child, py . Equation (1) describes the parent as a social planner who attaches weights 1 and α to himself and his child respectively. If the parent is altruistic toward the child, it must be case that

$$U(y) + \alpha U(c_2) \geq U(y) \text{ whenever } c_2 \geq 0.$$

This condition says that if parental consumption is the same whether the child is born or not, then the parent must be better off having the child. This condition is equivalent to the condition $U(c) \geq 0$.

Assume a CRRA utility function $U(c) = \frac{1}{1-\sigma} c^{1-\sigma} + \underline{C}$ where $\sigma \geq 0$. Constant $\underline{C} > 0$ guarantees positive utility when $\sigma > 1$. In this case consumption needs to be further restricted to be larger than $\omega := ((\sigma - 1) \underline{C})^{\frac{1}{1-\sigma}}$. The interpretation of ω is important: it corresponds to the imputed consumption level in the unborn state. This parameter has practical implications because altruistic

parents will never drive their child's consumption below ω .³ It is convenient to rewrite $U(c)$ as $U(c) = u(c) - u(\omega)$ where $u(c) = \frac{1}{1-\sigma}c^{1-\sigma}$, the standard CRRA representation.

The optimal allocation when a child is born satisfies $c_2 = c_1 (\alpha/p)^{1/\sigma}$ and $c_1 \left(1 + p (\alpha/p)^{1/\sigma}\right) = (1 - \lambda)y + py$. To further simplify suppose $p = \alpha$ so that

$$c_1 = c_2 = \frac{(1 - \lambda + \alpha)}{1 + \alpha}y$$

and $V_1 = U(c_1) (1 + \alpha)$. The child is born when $V_1 > V_0$, or

$$\frac{1}{1 - \sigma} (1 - \lambda + \alpha)^{1-\sigma} (1 + \alpha)^\sigma > \frac{1}{1 - \sigma} \left(1 + \alpha \left(\frac{\omega}{y}\right)^{1-\sigma}\right)$$

To better understand this condition, suppose first that $\omega = 0$ and $\sigma \in (0, 1)$. In this case the condition becomes $(1 - \lambda + \alpha)^{1-\sigma} (1 + \alpha)^\sigma > 1$. This restriction holds if λ is sufficiently low and/or α are sufficiently large. Crucially, the condition is always satisfied if σ is sufficiently close to one, for $\alpha > 0$ and $\lambda > 0$. The key insight of this result is that if altruistic parents are sufficiently averse to consumption dispersion, they would always have the child regardless of the cost. The reason is that a high σ implies strong diminishing returns to income, and having the child provides a way to escape the diminishing returns. This result is even stronger when $\sigma > 1$ and $\omega \rightarrow 0$.

A way to understand why parents prefer high fertility when σ is large is to compute the shadow price of a child, or the value of a child for short. This value can be defined as

$$VC = \frac{\alpha U(c_2)}{u'(c_1)}, \quad (2)$$

where the numerator is the utility the parent derives from the child's utility, and the denominator transforms this utility in consumption units.⁴ The equation above can be rewritten as:

$$VC = \frac{\alpha}{\sigma - 1} \left(\left(\frac{c_1}{\omega}\right)^{\sigma-1} - 1 \right) \cdot c_1 \quad (3)$$

Consider first the case $\omega = 0$ and $\sigma \in (0, 1)$. In that case $VC = \frac{\alpha}{1-\sigma} \cdot c_1$ so that the shadow price of a child increases proportionally with parental consumption. More importantly, the value of a child goes to infinite as σ approaches one, a result that also holds true when $\sigma > 1$ and $\omega \rightarrow 0$. This

³Although imputing a level of consumption in the "unborn state" is unfamiliar to many, it arises naturally in altruistic models with endogenous population because of the need of a full description of the consumption space. Welfare in the unborn state is analogous to welfare in a dead state arising in models of longevity. See, for example, Rosen (1988), Becker, Philipson and Soares (2005), Murphy and Topel (2006), Hall and Jones (2007), Jones and Klenow (2011). In those models ω is the imputed consumption in the dead state and individuals with consumption below ω would prefer to be dead. The welfare of the unborn also arises in normative models of endogenous population, as in Golosov, Jones and Tertilt (2007).

⁴This expression will naturally appear in our benchmark model below from the first-order condition with respect to the number of children.

result shows one key problematic implication of models with a low curvature of the utility function, as this leads to implausibly large values for a child’s life and therefore implausibly high fertility rates. This result can be avoid by properly selecting ω to match any desired value of a child, for any σ . This is in essence the procedure utilized by Murphy and Topel (2006) to match the value of statistical life for adults. This trade-off between the value of σ and the value of ω is illustrated in Table 1.

Table 1 presents alternative estimates of the value of child computed from equation (3). As we obtain below in our fully calibrated model, the value of α is set to 0.545, and the value of λ to 0.308. Lifetime income for a parent y is computed using an annual income of \$23,946, which corresponds to half of the median household income in the US.⁵ It is assumed the working lifespan is 40 years and the interest rate is 2%. As Table 1 shows, with $\omega \simeq 0$ the value of a child is very large for any $\sigma > 1$. For instance, for $\sigma = 1.25$, a value common in quantitative macroeconomics, the value of a child is around \$76 million dollars. In contrast, if $\sigma = 0.40$, the value of a child is \$485,000. A similar value of \$465,000 could be obtained with a $\sigma = 1.25$ but increasing the annualized ω to \$5,000. However, since the per capita poverty line is roughly \$23,000 for a family of four, this parameterization of ω would imply that no children with consumption below the poverty line would be born. This is problematic because the evidence suggests that on the contrary, poorer families tend to have more children. Table 1 parallels Table 2 in Murphy and Topel (2006). Theirs illustrates how the value of life varies as a function of the imputed consumption in the dead state and the EIS. Similar to the message in Table 1, in their exercise the value of life is increasing in σ and decreasing in ω .

As we discuss more extensively below, when fertility is endogenous the value of a child will be necessarily linked with the marginal cost of raising a child. The total cost of raising a child born in 2011 is around \$357,627 for a low-income family and \$496,093 for a middle-income family (details in the calibration section). What these figures suggests, together with the observation that poor children are in fact born, is that the value of σ consistent with the value of a child is at odds with the value of $\sigma > 1$ commonly used in quantitative macro. We show next that a model that disentangles the EIS from the EGS is able to reconcile the low EIS in the aggregate data ($\sigma > 1$) with the value of a child in models with endogenous fertility choice. The key is that the EGS is the one determining the value of a child.

2.2 Generalized preferences

In this section we formulate a framework that allows to define and distinguish the EIS from the EGS. For this purpose we extend the model of the previous section to include a life-cycle of length T and multiple children. In this case it is convenient to define the lifetime consumption of an

⁵Below we calibrate our full model using information on the cost of raising children for a typical US family with two parents and two children. We adjust the data so that in the calibrated model one parent will raise one child on average.

individual as a composite consumption C that takes the form

$$C = \left[\sum_{t=0}^T \beta^t c_t^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad \sigma > 0, \beta \in (0, 1). \quad (4)$$

Absent children, C is the only source of utility for an individual. The function defining composite consumption is a CES aggregator with elasticity of substitution $1/\sigma$ and weights β^t . As we show below, $1/\sigma$ is the standard EIS. Notice that $C \geq 0$ for all σ .

The lifetime utility of an individual, V , is described by the preferences

$$V \equiv \frac{1}{1-\eta} C^{1-\eta} + \int_0^n \varphi(i) V'_i di, \quad \eta \in (0, 1), \quad (5)$$

where n is the number of children, V'_i is the utility of child i and $\varphi(i) \geq 0$ is the weight that the parent attaches to child i . Positive weights means that the parent acts as a social planner at the family level where the implicit weight of the parent is 1. The key new parameter is η which determines the willingness to substitute composite consumption across generations. As we show below, $1/\eta$ is the elasticity of intergenerational substitution, EGS. Restriction $\eta \in (0, 1)$ is required for $V \geq 0$ and $V'_i \geq 0$, as otherwise adding a positive mass of children would be detrimental to parental utility.⁶ The utility described in the previous two-period example is a particular case of (5) which can be obtained by setting $\eta = \sigma$. In this case $V = \frac{1}{1-\sigma} \sum_{t=0}^T \beta^t c_t^{1-\sigma} + \int_0^n \varphi_i V'_i di$ which is the standard additively-separable formulation with a single elasticity.

We consider mostly the symmetric case $V'_i = V'$ for all i as is common in the literature. Symmetric treatment can arise as a way to avoid conflict among siblings. Furthermore, suppose $\varphi(i) = \alpha(1-\varepsilon)i^{-\varepsilon}$ which describes a particular type of discounting towards children: hyperbolic discounting.⁷ In this case $\int_0^n \varphi(i) di = \alpha n^{1-\varepsilon}$. Since parental weights are non-negative, function $\alpha n^{1-\varepsilon}$ must be positive and increasing which implies restrictions $\alpha > 0$ and $1 > \varepsilon > 0$. Symmetry and hyperbolic discounting simplifies (5) to

$$V = \frac{1}{1-\eta} C^{1-\eta} + \alpha n^{1-\varepsilon} V'. \quad (6)$$

This formulation is a generalized version of Becker and Barro (1988) that allows for $\eta \neq \sigma$. An alternative and convenient way to describe the same preferences is obtained by recursively substituting V' out of the equation. Under the boundedness condition $\lim_{t \rightarrow \infty} \alpha^t N_t^{1-\varepsilon} C_t^{1-\eta} / (1-\eta) = 0$, it follows that

$$V = \sum_{s=0}^{\infty} \alpha^s N_s^{1-\varepsilon} \frac{C_s^{1-\eta}}{1-\eta}, \quad (7)$$

where $n_0 = 1$ and $N_s = \prod_{v=0}^{s-1} n_v$. Notice that the summation in this equation is across generations,

⁶Section 5.2 extends the model to allow any $\eta > 0$.

⁷In Cordoba and Ripoll (2011) we discuss other types of discounting such as exponential discounting.

while the summation defining C_s in equation (4) is across ages, or over the life cycle. We now formally show that $1/\sigma$ and $1/\eta$ are the EIS and the EGS respectively.

The marginal rate of substitution between consumption at age s and consumption at age v for an individual is defined as $MRS(c_v, c_s) \equiv \frac{\partial V/\partial c_v}{\partial V/\partial c_s} = \frac{\partial C/\partial c_v}{\partial C/\partial c_s}$. The corresponding EIS between consumptions at period s and period v is then defined as

$$EIS(c_v, c_s) = \frac{d \ln(c_s/c_v)}{d \ln MRS(c_v, c_s)}. \quad (8)$$

The typical EIS in the literature refers to two adjacent periods, say s and $s+1$, but for the isoelastic preferences (6), the EIS is independent of s and v , as long as $s \neq v$. To see this, notice that

$$MRS(c_v, c_s) = \frac{\partial C/\partial c_v}{\partial C/\partial c_s} = \frac{C^{-\eta} C^\sigma \beta^v c_v^{-\sigma}}{C^{-\eta} C^\sigma \beta^s c_s^{-\sigma}} = \beta^{v-s} (c_s/c_v)^\sigma$$

and therefore $EIS = EIS(c_v, c_s) = 1/\sigma$.

The EGS can be defined similarly to the EIS but it relates to the willingness to substitute consumption across different generations rather than across different ages. The marginal rate of substitution between *composite* consumption of generations s and v from the point of view of the initial parent is given by $MRS(C_v, C_s) = \frac{\partial V/\partial C_v}{\partial V/\partial C_s}$. The corresponding EGS can be defined as

$$EGS(C_v, C_s) = \frac{d \ln(C_s/C_v)}{d \ln MRS(C_v, C_s)}. \quad (9)$$

According to this definition, the EGS measures the willingness of the parent to substitute composite consumption across generations s and v , $s \neq v$. Similarly to the EIS, the EGS could be in principle defined only for adjacent generations but, as Proposition 2 below shows, $EGS = 1/\eta$ for any $s \neq v$ when preferences are described by (7). An alternative definition of the EGS that does not involve composite consumption but specific consumptions for parents and children is

$$\widetilde{EGS}(c_v, c'_s) = \frac{\partial \ln(c'_s/c_v)}{\partial \ln(MRS(c_v, c'_s))}, \quad (10)$$

where c_v is parental consumption at age v and c'_s is children consumption at age s . The partial derivative refers to a change in the c'_s/c_v ratio holding the other consumption ratios constant. It turns out that $\widetilde{EGS}(c_v, c'_s) = \widetilde{EGS} = EGS = 1/\eta$ for any v and s' .

Proposition 2 - Isoelastic substitution. Suppose the lifetime utility of an individual is described by (6). Then $EIS(c_v, c_s) = EIS = 1/\sigma$ and $EGS(C_v, C_s) = \widetilde{EGS}(c_v, c'_s) = EGS = 1/\eta$.

Proof. The first part of the proposition was derived above. Next, using (7)

$$MRS(C_v, C_s) = \frac{\partial V / \partial C_v}{\partial V / \partial C_s} = \frac{\alpha^v N_v^{1-\varepsilon} C_v^{-\eta}}{\alpha^s N_s^{1-\varepsilon} C_s^{-\eta}},$$

and therefore

$$EGS(C_v, C_s) = \frac{d \ln(C_s / C_v)}{d \ln MRS(C_v, C_s)} = 1/\eta.$$

Moreover,

$$MRS(c_v, c'_s) = \frac{\partial V / \partial c_v}{\partial V / \partial c'_s} = \frac{C^{-\eta} C^\sigma \beta^v c_v^{-\sigma}}{\alpha n^{1-\varepsilon} (C')^{-\eta} (C')^\sigma \beta^s (c'_s)^{-\sigma}}.$$

Since C is constant returns to scale, it can be written as $C = c_v \widehat{C}_v$ where \widehat{C}_v is homogeneous of degree zero. As a result,

$$MRS(c_v, c'_s) = \frac{c_v^{-\eta} \widehat{C}_v^{-\eta} c_v^\sigma \widehat{C}_v^\sigma \beta^v c_v^{-\sigma}}{\alpha n^{1-\varepsilon} (c'_s)^{-\eta} (\widehat{C}'_s)^{-\eta} (c'_s)^\sigma (\widehat{C}'_s)^\sigma \beta^s (c'_s)^{-\sigma}} = \frac{\widehat{C}_v^{\sigma-\eta} \beta^v}{\alpha n^{1-\varepsilon} (\widehat{C}'_s)^{\sigma-\eta} \beta^s} \left(\frac{c_v}{c'_s} \right)^{-\eta}.$$

Consider a change in c_v/c'_s holding all other consumption ratios constant. In that case

$$\widetilde{EGS}(c_v, c'_s) = \frac{\partial \ln(c'_s / c_v)}{\partial \ln(MRS(c_v, c'_s))} = 1/\eta.$$

While estimates of the value of σ are available in the literature, much less is known about the value of both ε and η . Quantitative macro models typically use values of $\sigma > 1$ implying a low EIS in aggregate data. For instance, Hall and Jones (2007) use $\sigma = 2$ and Murphy and Topel (2006) set $\sigma = 1.25$. Parameter ε has not been directly estimated, but some models of parental altruism have calibrated it. For example, Birchenal and Soares (2009) calibrate values of ε in the range of 0.4 to 0.6, while Doepke (2004) calibrates $\varepsilon = 0.5$. Our new parameter η has never been estimated. Models in the tradition of Becker and Barro (1988) generally assume that $\sigma = \eta < 1$, so they do not have a conceptual distinction between the EIS and the EGS, and for technical reasons, they assume a value lower than one. Among this class of models, Doepke (2004) is perhaps the only one that calibrates $\sigma = \eta$ and obtains a value of 0.5. Different from this literature, we need a strategy to calibrate η as a different parameter from σ . As suggested in the previous section, the value of a child provides a way to identify the EGS.

We now link the value of η with the value of a child, which in a model with continuous n can be defined as a marginal rate of substitution between consumption and children. Specifically,

$$VC = \frac{\partial V / \partial n}{\partial V / \partial c_y}$$

where c_y is the consumption of the parent in period y , when the child is born. Using (6), the value

of a child is given by

$$VC = \frac{1}{1-\eta} \frac{(1-\varepsilon)\alpha n^{-\varepsilon} C^{1-\sigma}}{1-\alpha n^{1-\varepsilon} \beta c_y^{-\sigma}}. \quad (11)$$

This equation shows that it is η rather than σ , the parameter directly determining the value of a child. In other words, it is the EGS, not the EIS, what can be most directly identified from the value of the child. This stands in contrast with the role of σ in determining the value of a child in equation (3). Equation (11) makes clear that when the EIS and the EGS are disentangled, it is now possible to have $\sigma > 1$ as in quantitative macro models while at the same calibrating η to match the value of a child, as we show next.

3 The Model

In this section we fully specify an endogenous fertility model suitable for calibration. The model incorporates the preferences introduced in the previous section and allows for life-cycle and intergenerational savings. Generations are connected through parental transfers to children. We impose constraints to these intergenerational transfers by assuming children cannot borrow to cover their own expenditures during childhood, so they fully depend on parental resources during that period. As we extensively discuss in Cordoba and Ripoll (2014), these constraints are particularly important to model the link between fertility and parental income. More specifically, absent constraints to intergenerational transfers, if the child's future income is larger than the cost of raising him, even altruistic parents would have incentives to extract rents from the child. This unconstrained parent would borrow against the child's future income to cover the costs of raising him. By assuming that the child cannot borrow to cover his childhood expenses, we are precluding the parent from extracting any rents.

Consider an economy in which individuals live for three periods: one as a child, one as a young adult and one as an old adult. Young adults work and raise children, while children and old adults only consume. Let b be the total lifetime transfers from a parent to each of his children. The individuals' problem for $t \geq 0$ is:

$$V(b) = \max_{c_c, c_y, c_o, n, b'} \frac{1}{1-\eta} C(c_c, c_y, c_o)^{1-\eta} + \alpha n^{1-\varepsilon} V(b'), \quad (12)$$

subject to

$$\begin{aligned} Rb + w &\geq Rc_c + c_y + c_o/R + n(b' + \lambda w); \\ b &\geq c_c; \\ \bar{n} &\geq n \geq 0; \end{aligned} \quad (13)$$

where c_c , c_y , and c_o are the consumptions as child, young adult and old adult respectively, R is the gross interest rate, w is wage income, and λ is the time cost of raising a child. The first constraint

is the present value budget constraint of an individual: transfers from his parent plus labor income must cover consumption expenses plus the cost of raising children. The cost of a child is the time cost of the parent, λ , plus the amount of transfers per child, b' . The second constraint is the credit constraint: children cannot borrow and solely rely on parental transfers b to consume c_c until they become young adults and can work.

Composite consumption, C , is given by

$$C(c_c, c_y, c_o) = [c_c^{1-\sigma} + \beta c_y^{1-\sigma} + \beta^2 c_o^{1-\sigma}]^{\frac{1}{1-\sigma}} + \underline{C},$$

where $\underline{C} > 0$, non-market consumption, is a constant that allow for non-homothetic preferences. As discussed in Cordoba and Ripoll (2014), a form of non-homotheticity as well as constraints to intergenerational transfers allow deterministic dynastic altruistic models to replicate the observed negative relationship between fertility and income.

Some additional parametric assumptions are required. Similar to Becker and Barro (1988), the condition $\eta > \varepsilon$ is required to avoid zero children to be the optimal solution. In addition, and just for analytical simplicity, we assume $\beta R = 1$ so that, as shown below, $c_y = c_o$ and adults have a simple flat consumption when young and old. Our results are robust to alternative assumptions that give rise to a more realistic consumption profile over the life cycle. Finally, the following restriction ensures that childhood credit constraints bind in a steady state situation.

Assumption 2. $\beta > \lambda$.

To understand why this assumption guarantees a binding credit constraint, notice that in steady state the present value of the child's future income is $\beta w = w/R$, while λw is the time cost of raising a child. Then Assumption 2 states that children are a net financial gain. In such a situation, and absent any constraints to intergenerational transfers, parents would have the incentive to have as many children as possible in order to extract rents from them.⁸ However, if children cannot borrow and fully depend on parental resources during childhood, the parents' ability to extract rents is restricted. As a result, in steady state the childhood credit constraint binds and the parent transfers the child just enough to cover consumption, i.e., $b = b' = c_c$. In addition, the incentives to have as many children as possible disappear and the number of children is below the maximum, as we show below.⁹

⁸In Cordoba and Ripoll (2014) we show these results using a standard model with $\sigma = \eta$.

⁹In contrast with Assumption 2, it is useful to recall that one of the implications of the Becker and Barro (1988) model is that children are a net financial cost in the sense that the present value of the cost of raising a child is higher than the value of his future income. This is the case in that model because otherwise, if children were a net financial benefit, then parents would have the maximum possible number of them to maximize their rents. Underlying this implication in Becker and Barro (1988) is the fact that there are no constraints to intergenerational transfers, so parents could effectively borrow against children's future income. As we discuss in Cordoba and Ripoll (2014), if children are in fact a net financial benefit, but we do not see parents having the maximum number of children, then it must be that somehow parents cannot extract rents from their children. In fact, there are legal, moral or enforcement reasons why parents cannot do so. This motivates the introduction of our childhood credit constraint.

3.1 Optimal consumption and transfers

Let θ and μ be the Lagrange multipliers on the budget constraint and the credit constraint respectively. Optimal solutions are such that $b' > 0$ whenever $n > 0$. Otherwise children's consumption would be zero. Since $C(c_c, c_y, c_o)$ satisfies Inada conditions, zero consumption would not be optimal. The first order conditions with respect to c_c , c_y , c_o , and b' are:

$$C^{-\eta} (C - \underline{C})^\sigma c_c^{-\sigma} = R\theta + \mu; \quad (14)$$

$$C^{-\eta} (C - \underline{C})^\sigma \beta c_y^{-\sigma} = \theta; \quad (15)$$

$$C^{-\eta} (C - \underline{C})^\sigma \beta^2 c_o^{-\sigma} = \theta/R \text{ and} \quad (16)$$

$$\theta = \alpha n^{-\varepsilon} V_b(b'). \quad (17)$$

Furthermore, the envelope condition reads

$$V_b(b) = \theta R + \mu. \quad (18)$$

Equations (14), (15), (16), and the assumption $\beta R = 1$ imply that

$$c_y = c_o \geq c_c. \quad (19)$$

The last inequality is strict when the credit constraint is binding. Moreover, equations (17) and (18) can be written as

$$\theta \geq R\alpha n^{-\varepsilon} \theta'. \quad (20)$$

If the expression above holds with equality, then it becomes an intergenerational version of the standard Euler equation, with θ being the marginal utility of a young adult's consumption, and $\alpha n^{-\varepsilon}$ an endogenous discount factor, or average degree of altruism, which depends on the number of children.

In what follows we focus on a steady state situation. In this case (20) simplifies to $1 \geq R\alpha n^{-\varepsilon}$. If children could borrow then this expression would hold with equality and the steady state number of children would be $n^* = (R\alpha)^{1/\varepsilon}$. Fertility in the unconstrained case is thus a function of the interest rate but independent of any income or level variable such as wages or \underline{C} . This case would be inconsistent with the documented evidence of a negative relationship between fertility and income (see Jones and Tertilt, 2008). Alternatively, when the credit constraint binds, $c_c = b$, the strict inequality $1 > R\alpha n^{-\varepsilon}$ holds. This implies that steady state fertility in the constrained case is larger than in the unconstrained case. More precisely, combining (15), (17), (18), and (14), it follows that

$$C^{-\eta} (C - \underline{C})^\sigma \beta c_y^{-\sigma} = \alpha n^{-\varepsilon} [C'^{-\eta} (C' - \underline{C})^\sigma c_c'^{-\sigma}].$$

In steady state this equation simplifies to

$$c_y = G(n)^{1/\sigma} c_c, \quad (21)$$

where $G(n) \equiv n^\varepsilon/R\alpha \geq 1$. This expression is convenient because $G(n)$ equals 1 when credit constraints are not binding implying $c_c = c_y (= c_o)$. In contrast, when the constraint binds the magnitude of $G(n) > 1$ measures the "tightness" of the constraint or the extent to which consumption during childhood falls below that of the adult period, i.e., $c_c < c_y$. Furthermore, $G(n)$ is increasing in the number of children, implying that the larger the number of children, the tighter the credit constraint due to the fact that the average degree of altruism per child decreases with n . Thus, the constrained allocation captures a quality-quantity trade-off between the number of children and the resources parents spend on them during childhood. From now on we concentrate on the constrained allocation.

Given n , steady state solutions for consumptions and transfers can be obtained using (19), (21), (13) and the transfer constraint $b = c_c$ as

$$c_y = c_o = c(n) \equiv \frac{G(n)^{1/\sigma} w (1 - \lambda n)}{n + G(n)^{1/\sigma} (1 + \beta)}; \quad (22)$$

$$b = c_c = c_c(n) \equiv \frac{1 - \lambda n}{n + G(n)^{1/\sigma} (1 + \beta)} w; \quad (23a)$$

$$b + \lambda w = \frac{1 + G(n)^{1/\sigma} \lambda (1 + \beta)}{n + G(n)^{1/\sigma} (1 + \beta)} w. \quad (24)$$

According to these expressions, net transfers to a child b decrease with the number of children for two reasons: additional children lower the net income of parents as they reduce parental labor supply; and more children increase the discount per-child due to the decreasing degree of altruism. Furthermore, the total cost of a child, $b + \lambda w$, decreases with the number of children because parents reduce transfers per-child while the time cost per-child remains constant. Adult consumption, on the other hand may be decreasing or increasing in the number of children. A sufficient condition for adult consumption $c(n)$ to be decreasing in the number of children is $\sigma > \varepsilon$, which turns out to be the empirically relevant case because, as we argue in different parts of this paper $\sigma > \eta$ and $\eta > \varepsilon$ are needed for an interior solution for fertility.

Given solutions for consumptions, steady state C and V can be written as

$$C(n) = \left(G(n)^{1-1/\sigma} + \beta + \beta^2 \right)^{\frac{1}{1-\sigma}} c(n) + \underline{C} \quad \text{and} \quad V(n) = \frac{1}{1-\eta} \frac{C(n)^{1-\eta}}{1 - \alpha n^{1-\varepsilon}}. \quad (25)$$

Provided $\sigma > \varepsilon$, the utility the parent derives from own consumption, $C(n)$, is also decreasing in n .

3.2 Optimal fertility

We now turn to the fertility choice. The optimality condition for fertility in an interior solution is

$$w\lambda + b' = \varphi(n) \frac{V(b')}{\theta}, \quad (26)$$

where $\varphi(n) = (1 - \varepsilon) \alpha n^{-\varepsilon}$. The left-hand side of this expression is the marginal cost of a child, which includes the value of parental time cost plus all transfers. The right hand side is the marginal benefit of the n -th child. The term $V(b')/\theta$ is the welfare of the child measured in parental consumption units, while $\varphi(n)$ is the weight that parents give to the n child. It is convenient to define the marginal benefit as $VC(n) \equiv (1 - \varepsilon) \alpha n^{-\varepsilon} \frac{V(b')}{\theta}$ which, as we show next, corresponds to the concept of the value of a child discussed in Section 2. Using (25), in steady state $VC(n)$ can be written as

$$VC(n) = \frac{1}{1 - \eta} \frac{(1 - \varepsilon) \alpha n^{-\varepsilon}}{1 - \alpha n^{1-\varepsilon}} \frac{C}{(C - \underline{C})^\sigma \beta c(n)^{-\sigma}}$$

which exactly maps into equation (11) when $\underline{C} = 0$. As in Section 2, this equation highlights the role of η in determining the value of a child. Using (15) the expression simplifies to

$$VC(n) = \frac{1 - \varepsilon}{1 - \eta} \frac{G(n)^{1-1/\sigma} + \beta + \beta^2}{G(n) - n/R} \frac{c(n)}{1 - C/C(n)}. \quad (27)$$

This expression is a generalized version of (2) that includes adjustments for a life-cycle component, an infinite horizon for the dynasty, multiple children and diminishing altruism. A simple way to recover (2) for the case $\omega = 0$, is to let $n = \beta = 1$ and $\varepsilon = \underline{C} = 0$. In this case

$$VC(1) = \frac{\alpha}{1 - \eta} \frac{\alpha^{1/\sigma-1} + 2}{1 - \alpha} c(n)$$

which is almost identical to (2) except that here η is the relevant parameter, not σ ; term $(\alpha^{1/\sigma-1} + 2)$ reflects the effect of the binding transfer constraint; and the division by $(1 - \alpha)$ reflects the infinite horizon of the dynasty. Equation (27) makes clear the connection between the value of a child and the EGS.

Using (24), (22) and (27), the solution for steady state fertility is characterized by

$$\frac{w\lambda + b}{c(n)} = \frac{VC(n)}{c(n)},$$

or

$$\frac{1 + G(n)^{1/\sigma} \lambda (1 + \beta)}{1 - \lambda n} = \frac{1 - \varepsilon}{1 - \eta} \frac{G(n) + G(n)^{1/\sigma} \beta (1 + \beta)}{G(n) - n/R} \frac{1}{1 - \underline{C}/C(n)},$$

which equates the marginal cost and marginal benefit of a child, both as proportion of parental consumption. In order for the credit constraint to bind, it must be the case the marginal benefit

is larger than the marginal cost at $n = n^*$. Notice that $G(n^*) = 1$ and therefore the following parameter restriction is required (assuming $\underline{C} = 0$).

$$\frac{1 + \lambda(1 + \beta)}{1 + \beta(1 + \beta)} \frac{1 - \beta(R\alpha)^{1/\varepsilon}}{1 - \lambda(R\alpha)^{1/\varepsilon}} < \frac{1 - \varepsilon}{1 - \eta}$$

Since $\eta > \varepsilon$, then this condition is satisfied when $\lambda < \beta$, which explains the need for Assumption 2. This turns out to be the case in the calibration.

4 Calibration and results

In this section we calibrate the model and provide a way to identify the EGS as a parameter distinct from the EIS.

4.1 Fertility data

We use data from a cross-section of US states to calibrate the model. This data is optimal for our purpose for several reasons. First, in contrast with cross-sectional international data in which countries are at different stages of the demographic transition, US states have all completed this transition. This feature maps better into our steady-state analysis. Second, cross-state data is better for our purpose than individual-level data because relative income across states is roughly constant, while individual income in any given year does not represent lifetime income. In this respect, the cross-state fertility-income relationship is closer to the one captured in the model. Third, despite the relative convergence in both income and fertility across US states, there is still some cross-sectional variation. Last, the assumption that the interest rate is identical can be better justified across US states than across countries.

Table 2 reports descriptive statistics for the total fertility rate in 2010 and median household income across US states, as well as their values for a subsample. The total fertility rate is from the 2012 National Vital Statistical Report, and it corresponds to the number of births 1,000 women age 15-44 would have in their lifetime if they experienced the births currently occurring at each age. Median household income is from the Statistical Abstract of the US (Census Bureau, 2012). We use the average median household income 2004-2006 to exclude the recent recession. Average total fertility in the sample is 1.944 children with a standard deviation of 0.175, while average median household income is \$47,892 with a standard deviation of \$7,178. As the table shows, relatively poorer states like Arkansas, Oklahoma and Texas have higher than average total fertility. In contrast, states with higher-than-average income such as Massachusetts, New York and Rhode Island have below-average fertility. Utah is one of the exceptions, with high income and high fertility. Although an outlier, Utah is a relatively small state in terms of population.

Figure 1 plots the total fertility rate versus median household income. The size of the bubbles represents 2010 population weights from the Statistical Abstract of the US. Taking into account

population weights, Figure 1 suggests a slightly negative relationship between fertility and income. Base on this data we estimate an income elasticity of fertility of -0.143 (significant at the 5% level on a population-weighted regression). This elasticity is close to the one estimated by Jones and Tertilt (2008) using individual-level Census data for the most recent cohorts. For instance, for the 1951-1955 cohort, whose average fertility was 2.05 children ever born and average occupational income was \$49,378, they estimate an income elasticity of fertility of -0.17 . We will use our estimated elasticity as one of the calibration targets next section.

4.2 The costs of raising children

The costs of raising children are fundamental for the calibration of our model. Our calibration requires data on both the goods costs and the time costs of raising children. Recall that in our model $b = c_c$ effectively corresponds to the present value of the goods costs of raising a child, while λw is the present value of the time costs. We use data from the USDA (2012) to compute the goods cost of raising a child. According to the USDA (2012), the typical cost of raising a child born in 2011 from age 0 to 17 for a family of four in the lowest income group is \$169,080, while for a family in the middle-income group is \$234,900 and for a high-income family is \$389,670 in 2011 dollars. These figures include direct parental expenses made on children through age 17 such as housing, food, transportation, health care, clothing, child care, and private expenses in education. Costs are projected using an inflation rate of 2.55%. Assuming a discount rate of 2%, the corresponding present values of these sums are \$143,051 for low income, \$198,437 for middle income, and \$328,990 for the high income group.

Table 3 presents these goods costs for a "representative family" in each income group. Using the family income brackets from USDA (2012), we select a 2011 income of \$43,625 for the representative low-income family; \$81,140 for middle income, and \$126,435 for a high-income family. The low-income family figure is computed as the average of the following two values: \$27,840, which corresponds to the income of a family in which both parents make the federal minimum wage in 2011; and \$59,410, which is the upper bound of low-income families from the USDA (2012) classification. The middle-income family number is simply the mid-point of the USDA (2012) interval of \$59,410 to \$102,280. Last, the high-income "representative" family is computed as the average between \$102,280 and \$150,000, where the latter corresponds to the 90th percentile of the family income distribution in 2011 according to the US Census Bureau. For each of the representative families we also compute a lifetime household income assuming a 40-year working life span and a 2% real interest rate. Comparing Table 2 and Table 3, notice that the values of median household income by states fall in between the low and middle-income family groups under the USDA classification. As a result, we will use the information of these two groups to calibrate the goods costs of raising children.

As we discuss in Cordoba and Ripoll (2014), accounting for the time costs of raising children is not trivial. Available estimates are based on time use survey data, but the difficulty of measuring time costs is that in many instances parents multitask, taking care of children as a secondary

activity while performing other primary activities. Using the 2003-2006 American Time Use Survey, Guryan, Hurst and Kearney (2008) find that while mothers spend around 14 hours per week in child care, fathers spend around 7 hours. These measures only include the time parents spend primarily on basic care of children, education, recreation and any travel related to these. They refer to overall averages for families with at least one child under the age of 18. However, if the total time parents spend in the presence of their children is measured (both primary and secondary time), then mothers spend 45 hours per week and fathers spend 30 hours. The extent to which both primary and secondary time should be included in the cost of raising children is a matter of debate in the literature.

In a related study, Folbre (2008) uses the 1997 Child Development Supplement of the Panel Survey of Income Dynamics to conclude that the average amount of both passive and active parental-care hours per child (not including sleep) is 41.3 per week for a two-parent household with two children ages 0 to 11. Passive care corresponds to the time the child is awake but not engaged in activity with an adult, while active parental care measures the time the child is engaged in activity with at least a parent. In addition to reporting hours spent in child care, Folbre (2008) discusses two alternative ways of computing the monetary value of these hours: one uses a child-care worker's wage and the other the median wage. When the former method is used in combination with the USDA (2012) goods cost of raising children, the time cost of raising children is on average around 60% of the total costs (see Table 7.3, p. 135), a lower-end estimate. Since the median wage is around the double of a child-care worker's wage, then using the former time valuation the time cost of raising children increases to 75% of the total costs. This evidence suggests the time costs of raising children are high: they are between 1.5 and 3 times the goods costs or direct expenditures in children.

In order to compute the time costs for each of the representative families in Table 3 we use the more conservative estimate in which they are about 60% of the total cost of raising a child. This more conservative estimate holds either when the 21 hours per week of primary care in Guryan, Hurst and Kearney (2008) are valued at the median wage, or when the 41.3 hours per week in Folbre (2008) are valued at a child-care worker's wage. As can be seen in Table 3, the present value of the time costs of raising a child is \$214,576 for a low-income family, \$297,656 for a middle-income family, and \$493,485 for a high-income family. Table 3 also presents the total costs of raising a child: \$357,627 for a low-income family, \$496,093 for a middle-income family, and \$822,476 for a high-income family.

It is important to notice that the total costs of raising a child in Table 3 map into the value of a child. As discussed before, optimal fertility is decided comparing the marginal cost and the marginal benefit of a child. Since the values in Table 3 correspond to the total marginal cost of a child, they also correspond to the value of a child, or the marginal benefit. As the median household income across US states in Table 2 ranges from \$35,261 to \$64,169, our calibrated model should be consistent with a value of child ranging between \$357,627 and \$496,093, which correspond to the total cost of raising a child for low and middle-income families in Table 3.

There is a parallel literature that estimates the "statistical" value of a child. Birchenal and

Soares (2009) survey this literature and report a range of estimates between \$1.3 and \$4 million. These estimates are taken from studies that consider purchases of safer cars, car seats and bicycle helmets to reduce the risk of death of a child. The upper bound of \$4 million is an order of magnitude higher than the range of the value of a child we will use in our calibration, between \$357,627 and \$496,093. Our model does not include mortality risk, so it cannot be calibrated to the "statistical" value of a child. However, as Birchenal and Soares (2009) point out, the statistical value of a child can be rationalized if the utility of the parent is modified to include an emotional loss associated to the death of a child. Absent this emotional loss, the value of a child must equal the marginal cost of raising a child in an interior solution.

A useful way of presenting the information in Table 3 for calibration purposes is to compute the costs of raising a child as a fraction of the lifetime household family income. These figures are presented in Table 4. As the table indicates, for the average two-parent two-child family in the low and middle-income groups in the US, the goods costs of raising each child correspond to 10.3% of the lifetime household income, while time costs are 15.4% and total costs are 25.6%. As we now turn to discuss, these figures will be key in our calibration.

A remark regarding parental transfers in our model is in order. Our model implies that when transfer constraints bind, no bequests or other inter-vivos transfers to adult children are given. All the transfers the parent gives to the child are spent during the childhood period ($b = c_c$). As we discuss in Cordoba and Ripoll (2014), although inter-vivos transfers and voluntary bequests do occur in the United States, a relatively small fraction of adults receive them, and they occur in small amounts. For instance, using the 1988 special supplement on transfers between relatives from the PSID, Altonji, Hayashi and Kotlikoff (1997) document that only 23% of adult children (on average 31 years old) receive transfers from parents (on average 59 years old). These are overall small transfers: the mean is \$3,442, and the median is \$951 in 2011 dollars. A similar pattern has been documented for bequests. Using the 1993-1995 Asset and Health Dynamics among the Oldest Old (AHEAD) data, Hurd and Smith (2001) document that most bequests are of little or no value: single descendants at the bottom 30% receive \$2,952, and the average single descendant receives \$14,760 in 2011 dollars. Given the highly skewed wealth distribution in the United States, the occurrence of significant bequests concerns only of a small fraction of the population.

4.3 Calibrated parameters

Some parameters in the model are set exogenously. We set the length of each of the three periods of life to 25 years: a child consumes with the resources transferred by his parent from ages 0 to 25; young adults have children at age 25 and work until age 50, while old adults consume and retire from age 50 to 75. The annual interest rate is set to 2% which implies a discount factor β of 0.61 per 25-year period. Finally, we set $\sigma = 1.5$, a value commonly used in quantitative macro models. This value implies that $EIS = 0.67 < 1$, a low rate of intertemporal substitution.

The remaining five parameters $[\lambda, \alpha, \varepsilon, \eta, \underline{C}]$ are calibrated to five targets. Although the model is non-linear and these parameters are jointly calibrated, each parameter can be more directly related

to one of the targets. Table 5 presents the results of our calibration exercise. First, parameter λ corresponds to the present value of the time costs of raising a child, λw , as a fraction of parental lifetime income w . We calibrate λ to match this share in the data, which according to Table 4 corresponds to an average of 15.4% for a typical family with two parents and two children in the low and middle-income USDA (2012) groups. Since in our model there is a single parent, and the average fertility in the sample is two children per household, then the average single parent will be raising one child and $\lambda = 0.308$.

Second, parameter α , which corresponds to the level parameter in the utility weight the parent gives to the children, has a first-order effect on the ratio b/w . This corresponds to the transfers the parent gives to each child relative to parental lifetime income. To see this, using equation (23a) write

$$\frac{b}{w} = \frac{1 - \lambda n}{n + \alpha^{-1/\sigma} \beta^{1/\sigma} n^{\varepsilon/\sigma} (1 + \beta)}.$$

Given the exogenous values of R , β and σ , as well as the calibrated value of λ , and given that as we discuss below the average n will be calibrated to a target of $n = 0.972$, then the equation above determines α for a given b/w target. According to Table 4, the average b/w for low and middle-income families is 10.3%. We obtain a calibrated value $\alpha = 0.545$.

Parameter ε determines the degree of diminishing altruism. To the best of our knowledge this parameter has not been directly estimated in the literature. Estimating this parameter requires to know how the parental valuation of children's utility falls as the number of children increases. The best way to back out the value of ε we could find is from a study by Dickie and Messman (2004). This paper directly uses stated-preference data on parental willingness to pay to relieve symptoms in children's acute respiratory illnesses. More importantly, the distinct feature of this study is that is the only one that estimates how parental willingness to pay changes with the number of children in the family. In addition to strongly supporting parental altruism toward their children, the paper estimates an elasticity of the parental willingness to pay with respect to the number of children in the family of -0.288 (see Table 5, p. 1159).

In order to map this elasticity into our model, and to the extent that health expenditures in treating acute illnesses increases the survival probability of the child, we compute the implied willingness to pay WTP for an increase $\Delta\pi_c$ in survival. It turns out that in this case the WTP

is directly linked to the value of a child and given by ¹⁰

$$WTP(n) = VC(n) \cdot \Delta\pi_c.$$

Since we are only interested in the elasticity of $WTP(n)$ with respect to n , the magnitude of term $\Delta\pi_c$ does not play a role in the value of this elasticity. Using the expression above together with equation (27) we obtain

$$\varepsilon^{WTP}(n) = \frac{\partial WTP}{\partial n} \frac{n}{WTP} = \frac{(1 - \varepsilon) - 1 + \alpha n^{1-\varepsilon}}{1 - \alpha n^{1-\varepsilon}}.$$

Given the calibrated α , and a calibration target of average $n = 0.972$, then the equation above determines ε . As Table 5 indicates, we obtain $\varepsilon = 0.676$. This value of ε is comparable to the corresponding parameter calibrated in Birchenal and Soares (2009). They obtain a value of 0.605 for their lower-bound estimate of the statistical value of a child (see their Table 2, p. 292).

Fourth comes our most important parameter, η . Although η is calibrated jointly with the rest of the parameters, its value is mainly identified from the value of a child. In particular, we choose η so that the model delivers an average fertility of around one child per parent, more precisely $n = 0.972$, which corresponds to half of the average fertility across US states on Table 2. Using the optimality condition for fertility in equation (26) together with the expressions for the value of a child in (27), and adult consumption (22), we can write

$$\frac{w\lambda + b}{w} = \frac{VC(n)}{w} = \frac{1 - \varepsilon}{1 - \eta} \frac{G(n)^{1-1/\sigma} + \beta + \beta^2}{G(n) - n/R} \frac{G(n)^{1/\sigma} (1 - \lambda n)}{n + G(n)^{1/\sigma} (1 + \beta)} \frac{1}{1 - \underline{C}/C(n)},$$

where the left-hand side represents the total cost of raising a child as a share of parental lifetime income, and the right-hand side is the value of a child also as a fraction of parental lifetime income. For $\underline{C} = 0$, and given the calibration targets described for λ , α , and ε , the equation above identifies η for a target of $n = 0.972$. Parameter \underline{C} is still to be determined, but as long as $\underline{C}/C(n)$ is small,

¹⁰To derive the $WTP(n)$ formula we write the utility of the parent as

$$V(b) = \max \frac{1}{1 - \eta} U(c_c, c_y, c_o)^{1-\eta} + \alpha(\pi_c n)^{1-\varepsilon} V(b'),$$

where π_c is the survival probability of the child. This implies that the willingness to pay to increase survival by $\Delta\pi_c$ is given by

$$WTP(n) = \frac{\partial V / \partial \pi_c}{\partial V / \partial c_y} \frac{\Delta\pi_c}{n}.$$

Because π_c and n enter symmetrically in the altruistic weight function, it turns out that

$$WTP(n) = \frac{\partial V / \partial \pi_c}{\partial V / \partial c_y} \frac{\Delta\pi_c}{n} = \frac{\partial V / \partial n}{\partial V / \partial c_y} \frac{\Delta\pi_c}{\pi_c} = VC(n) \cdot \Delta\pi_c.$$

Since in our model $\pi_c = 1$, evaluating the expression above at $\pi_c = 1$ delivers

$$WTP(n) = VC(n) \cdot \Delta\pi_c.$$

η will be of first-order importance in determining the value of a child, and through this channel, the fertility level. We obtain a calibrated value of $\eta = 0.35$.

Last is parameter \underline{C} , non-market consumption. We calibrate this parameter in order to match the income elasticity of fertility in our sample, which we computed to be -0.143 as described before. If $\underline{C} = 0$ then fertility would not be related to income in our model. To see that write the optimality condition of fertility in (26) as

$$w\lambda + b = \frac{1 - \varepsilon}{1 - \eta} \frac{G(n)^{1-1/\sigma} + \beta + \beta^2}{G(n) - n/R} \frac{c(n)}{1 - \underline{C}/C(n)}.$$

Since from (24) and (22) it follows that both the total cost of raising a child $w\lambda + b$ and adult consumption $c(n)$ are proportional to w , if $\underline{C} = 0$, then both the marginal cost and the marginal benefit of a child are proportional to lifetime income and fertility choices would be independent of w . Only when $\underline{C} > 0$ we have a link between fertility and income. In fact, this relationship is negative because when w increases the marginal benefit increases less than the marginal cost due to the presence of $\underline{C}/C(n)$. Calibrating \underline{C} in order to target an income elasticity of fertility of -0.143 results in a maximum $\underline{C}/C(n)$ of 5.73% across US states, a small value.¹¹

4.4 Robustness

We check the robustness of our calibration, especially the obtained value of $\eta = 0.35$, by performing two alternative exercises. We present these on Table 6. First, under alternative # 1, we select different target values for the goods and time costs of raising children. Going back to Table 4, rather than using the average between low and middle-income family groups we use the statistics for the low-income family group under the USDA (2012) classification. After all, the upper bound of household income for the low-income USDA group is \$59,410, and only four states in our sample have median household income above this level. Under this alternative calibration the target for the goods cost of raising a child as a share of household lifetime income would be 11.8%. The corresponding target for the time costs would be 17.6%. The rests of the target values remain the same as in Table 5. As Table 6 indicates, under this alternative calibration # 1 we obtain $\eta = 0.41$. The reason is that the new targets imply a larger value of a child, which require a larger value of η , or a lower EGS. Under alternative # 2 in Table 6, we set $\underline{C} = 0$ so that fertility is independent of income. In this case we obtain $\eta = 0.359$, a value close to our benchmark calibration. This is consistent with the fact that in our benchmark calibration \underline{C}/C is small. Overall, our alternative calibrations support $\eta < 1$, namely an EGS lower than one.

¹¹Although the income elasticity of fertility we estimate is statistically significant, its absolute value is small. Since in the calibrated model US states only differ in household income, this small elasticity implies that the model cannot be expected to explain all the fertility dispersion in the data. However, the important thing to notice here is that this does not affect the main point of this calibration exercise, which is to illustrate how our main parameter η can be identified from the value of a child that matches average fertility.

4.5 Discussion

Our calibrated model has implications for cross-state variations in the value of a child. Table 7 illustrates this point. Under our benchmark calibration, the average value of a child across US states is \$373,663. The maximum value of a child in the sample is \$509,572 and the minimum is \$268,251. One should bear in mind that our calibration is based on median household income by state, which given US inequality maps into the low and middle-income family groups under the USDA classification. Although Table 7 also illustrates the dispersion on the value of a child across US states, notice that the value of a child as a share of lifetime household income is similar across states, around 27.9% on average. According to Table 4, this number in the data is about 25.6%. More in general, our model could rationalize the higher economic value of child for higher-income families. For instance, Table 3 shows how for a representative household with annual income of \$126,435, the value of child is estimated to \$822,476. Given that the costs of raising children are proportional to family income in the data, and that in the model there is a link between the total cost of raising children and the value of a child, our model could also replicate this higher value.

The most important finding of this calibration exercise is that the EIS is very different from the EGS. While $\sigma = 1.5$ is a standard value in quantitative macro models, we obtain a calibrated value of $\eta = 0.35 < \sigma$. Our calibration suggest that individuals have a low intertemporal substitution, but a higher intergenerational substitution, a novel result. One of the central conclusions of this paper is that while the EIS mostly influences shorter-term economic decisions within the life span of an individual, the EGS mostly influences longer-term economic choices, those who affect more than one generation.

The EGS is conceptually and quantitatively different from a long-term EIS. For example, Biedermaier and Goenner (2008) allow the degree of intergenerational substitution to vary over the life cycle, so that short-term and long-term EISs emerge. They find that the EIS seems to be smaller and below one for longer time horizons. Conceptually, however, the long-term EIS still refers to an intertemporal willingness to substitute consumption across time for the same individual, and therefore it is different from the EGS, which refers to different individuals.

What would happen in our model if $\sigma = \eta = 0.35$? It turns out that our model could still be calibrated to match the same set targets, whether $\sigma = 1.5$ or $\sigma = 0.35$. However, such a calibration poses a number of important problems. The first one is that $\sigma = 0.35$ conflicts with extensive evidence based on aggregate consumption data, which supports an EIS lower than one. As discussed in Guvenen (2006), the largest EIS that has been either estimated econometrically or calibrated in the context of a model is at most one. In fact, when it comes to aggregate consumption data, the EIS is lower than one because it reflects the preferences of poorer individuals, those who do not participate in financial markets.

Second, a value of $\sigma < 1$ runs into problems when we consider the models at the intersection of macroeconomics and demographics. For instance, a number of macro papers examining life expectancy and health find that $\sigma > 1$ is necessary for the model to be consistent with facts. First, Hall and Jones (2007) analyze the raising share of income devoted to health spending in the US economy and show that the restriction $\sigma > 1$ is required to explain why longevity is a superior

good. In their paper, a low EIS implies a strongly diminishing marginal utility of income, while the marginal benefit of life extensions remains bounded. This feature of preferences explains why richer individuals want to spend an increasing fraction of their income in health in order to prolong their life span. Second, Jones and Schoonbroodt (2010) find that a low EIS, of about one third, is required to explain the fall of overall birth rates in response to falling infant mortality rates during the US demographic transition. In sum, the available literature suggests that models at the intersection between macroeconomics and demographics favor $\sigma > 1$, or a low EIS.

Our model does not consider mortality risk. However, if we introduced it, a calibration with $\sigma = \eta < 1$ would have the issue that it would not be consistent with the available evidence on the value of statistical life (VSL). Specifically, it can be shown that parameter σ plays a crucial role in determining the VSL, and that the VSL is increasing in σ . If $\sigma = \eta = 0.35$ the VSL would be much lower than the typical range of \$4.5 to \$9 million reported by Viscusi and Aldy (2003) for the US. In fact, the VSL implied by such calibration would be about \$1.5 million. In contrast, our calibration with $\sigma = 1.5$ and $\eta = 0.35$ is consistent a VSL of \$4.5 million.¹² In sum, our model with $\sigma > 1$ and $\eta < 1$ is better suited to describe demographic facts related to both fertility and mortality in the context of quantitative macro models.

5 Extensions

The main purpose of this section is to illustrate the scope of our framework beyond altruistic models of fertility choice. We provide a few extensions in order to illustrate other contexts in which disentangling the EGS from the EIS is useful. In particular, we suggest how to extend our framework to include infant mortality risk, to allow for $\eta > 1$, and to analyze long-term inequality in a model with idiosyncratic risk.

5.1 EGS and the coefficient of risk aversion

Those familiar with the Epstein-Zin-Weil (EZW) preferences from Epstein and Zin (1989) and Weil (1990), may find a resemblance between these and our formulation, and may wonder whether our framework reduces to a relabeling of EZW preferences. Although non-separability is a feature of both EZW and our preferences, they are conceptually quite different. EZW preferences disentangle aversion to risk from aversion to deterministic consumption fluctuations. In the absence of risk, the EZW formulation collapses into the standard formulation. This is not the case with our preferences. The framework we presented above does not model risk. Our preferences disentangle aversion to two types of *deterministic* fluctuations in consumption: (i) temporal variations; and (ii) intergenerational variation.

In order to illustrate the relationship between EZW and our preferences, we now introduce child mortality risk into our model, and combine EZW with our approach to disentangle three parameters: the EIS, the EGS and the coefficient of relative risk aversion (CRRA). Infant mortality

¹²Details of the calibrated version of our model with mortality risk are available upon request.

is a potentially important determinant of fertility choices.¹³ In order to introduce risk, it is convenient to utilize the following monotonic transformation of our preferences to recover the EZW preferences. Defining $W = [(1 - \eta)V]^{\frac{1}{1-\eta}}$, equation (5) can be rewritten as

$$W = \left[C^{1-\eta} + \int_0^n \varphi(i) W_i^{1-\eta} di \right]^{\frac{1}{1-\eta}}. \quad (28)$$

Notice that while V can be negative, W is non-negative so that zero is a lower bound, a property that we use shortly. Consider now the possibility that the lifetime utility of the child is a random variable, \widetilde{W} . Let $\mu(\widetilde{W})$ denote the certainty equivalent operator. In particular, Epstein and Zin (1989) as well as Weil (1990) consider a particular CRRA operator $\mu(\widetilde{W}) = [E\widetilde{W}^{1-\rho}]^{1/(1-\rho)}$ where $\rho \geq 0$ is the coefficient of relative risk aversion. For example, $\rho = 0$ means that parents are neutral to risks associated to their children's welfare. Following EZW, when certainty equivalent $\mu(\widetilde{W}_i)$ is what the parent perceives as the utility of his child, preferences can be described by

$$W = \left[C^{1-\eta} + \int_0^n \varphi(i) \mu(\widetilde{W}_i)^{1-\eta} di \right]^{\frac{1}{1-\eta}}, \quad (29)$$

Suppose now infant mortality is the only risk. In particular, let π be the survival probability of a newborn. In that case, $\mu(\widetilde{W}_i) = [\pi W_i'^{1-\rho} + (1-\pi)D]^{1/(1-\rho)}$ where D is the utility in case of death, or perhaps better, the perceived utility. To simplify, suppose $D = 0$ which means that being alive is always better than not, $W_i' \geq D = 0$. Furthermore, if the death of a child is not so painful as to eliminate all enjoyment of having children, then the additional assumption $\rho \in (0, 1)$ is required. In other words, if $\rho > 1$ so that parents are significantly risk averse, then $\mu(\widetilde{W}_i)$ would be zero whenever $D = 0$. Finally, assuming symmetric treatment of children, $W_i' = W'$, parental welfare simplifies to

$$W = \left[C^{1-\eta} + \alpha n^{1-\varepsilon} \pi^{(1-\eta)/(1-\rho)} W'^{1-\eta} \right]^{1/(1-\eta)}.$$

In order to relate the expression above to our earlier formulation in (6), it is convenient to rewrite preferences in terms of V rather than W to obtain

$$V = \frac{1}{1-\eta} \left(\sum_{t=0}^T \beta^t c_t^{1-\sigma} \right)^{\frac{1-\eta}{1-\sigma}} + \alpha n^{1-\varepsilon} \pi^{(1-\eta)/(1-\rho)} V'. \quad (30)$$

These preferences are an extension of our framework that disentangles three different concepts: the $EIS = 1/\sigma$, $EGS = 1/\eta$ and the $CRRA = \rho$. The expected utility model is the special case $\eta = \rho$, while if $\sigma = \eta = \rho$ would imply additive separability across time, generations, and states. Finally, if $1 - \varepsilon = \frac{1-\eta}{1-\rho}$ then parents only care about the number of surviving children, πn , which provides

¹³See, among others, Doepke (2005) and Jones and Schoonbroodt (2010).

microfoundations to the simplifying assumption made in the literature.¹⁴

5.2 EGS less than one

Our benchmark formulation assumes $\eta \in (0, 1)$, and the calibration shows this assumption is not binding. We now show that it is simple to relax this assumption. Consider the preferences in (28). As mentioned above, they are a monotonic transformation of our benchmark preferences and are strictly non-negative for any η , not just for $\eta \in (0, 1)$. In spite of W being positive, it is still true that parental welfare is decreasing in the number of children, n , when $\eta > 1$. Therefore, in that case the optimal number of children would be zero. This result, however, is due to the implicit assumption that the welfare of the unborn individual is zero (see Cordoba and Ripoll, 2011). The following generalized version of (28) makes this point clear. Suppose there is a number of potential children, n_p . Let W_i be the welfare of an individual if born, and D_i if unborn. D_i is what parents perceive, or impute, is the welfare of the unborn. This is analogous, but not the same, to the perceived utility in case of dead. As in the previous example, D could be normalized to zero so that altruistic parents perceive potential children are better off being born than unborn.¹⁵ In this case parental preferences are

$$W = \left[C^{1-\eta} + \int_0^n \varphi(i) W_i^{1-\eta} di + \int_n^{n_p} \varphi(i) D_i^{1-\eta} di \right]^{\frac{1}{1-\eta}}. \quad (31)$$

Equation (28) is a special case of (31) that requires $D = 0$. Notice that if $\eta > 1$ and $D = 0$ then $W = 0$. In this case, the small degree of substitutability between utilities means that if one individual receives zero utility then parents utility is also zero. To make an analogy with the theory of the firm, if W is production and the inputs are the utilities of individuals, then $\eta > 1$ means that all inputs are essential. To avoid this implication when assuming $D = 0$ requires the restriction $\eta \in (0, 1)$ as in the benchmark. But if $D > 0$ then parental utility increases with the number of children, for any $\eta > 0$, as long as $W \geq D$. Allowing for $D = 0$ and calibrating the model with preferences (31) would still require a low D and $\eta \in (0, 1)$, as otherwise the model would not be able to match the value of a child as discussed in Section 2.1.

5.3 EGS in Bewley models

The EGS is a potentially important determinant of long run inequality. To illustrate this point, consider an intergenerational version of a Bewley model extended to disentangle the EGS from the EIS. The source of inequality is uninsurable idiosyncratic risk in earning ability.

The economy is populated by a continuum of individuals of mass one who differ in their earning abilities, ω , and the amount of transfers they received from their parents, b . All individuals work from ages s to S and have one child at age F . For a given amount of transfers to his child, b' , the

¹⁴See, for instance, Jones and Schoonbroodt (2010).

¹⁵One way, although not the only way, to rationalize abortion by altruistic parents would occur when $D > W$.

individual allocates lifetime consumption by solving the problem

$$C = \max_{c_t} \left[\sum_{t=0}^T \beta^t c_t^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \quad \text{subject to} \quad \sum_{t=0}^T \frac{c_t}{R^t} = b + \sum_{t=s}^S \frac{\omega}{R^t} - \frac{b'}{R^F}.$$

Assuming $R = 1/\beta$, the optimal solution is

$$c_t = c_0 = \frac{1-\beta}{1-\beta^{T+1}} \left(b + \omega \frac{1-\beta^{S-s+1}}{1-\beta} - \beta^F b' \right) \quad \text{for } t = 0, \dots, T,$$

and

$$C = \left(\frac{1-\beta^{T+1}}{1-\beta} \right)^{\frac{\sigma}{\sigma-1}} \left(b + \omega \frac{1-\beta^{S-s+1}}{1-\beta} - \beta^F b' \right). \quad (32)$$

The ability of the child, ω' , is random and drawn from the distribution $\omega' \sim F(\omega'|\omega)$. Optimal transfers, b' , solve the problem:

$$V(b; \omega) = \max_{b' \geq 0} \frac{C^{1-\eta}}{1-\eta} + \alpha E [V(b'; \omega')|\omega] \quad \text{subject to (32)}.$$

Let $b' = g(b, \omega)$ be the optimal transfer rule and $p_0(b, \omega)$ the initial distribution of transfers and abilities. The evolution of the transfers-ability distribution, $p_{t+1}(b', \omega')$, can be calculated as:

$$p_{t+1}(b', \omega') = \sum_{\omega} \sum_{\{b: b'=g(b, \omega)\}} p_t(b, \omega) F(\omega'|\omega).$$

Let $p(b, \omega)$ be the associated invariant distribution, and $p(b) = \sum_{\omega} p(b', \omega')$ the distribution of lifetime transfers.

The model just described is an intergenerational Bewley model. Two constraints are important. One is that parents cannot insure their children against their random abilities because transfers are non-contingent. Second, transfers cannot be negative. We now calibrate two versions of the model to investigate the implications of disentangling the EGS from EIS. The first version uses the benchmark parameters of our calibration $[\beta, \alpha, \eta, \sigma] = [0.61, 0.545, 0.35, 1.5]$. For abilities, we assume they follow an AR(1) process $\ln \omega' = \rho \ln \omega + \epsilon$, $\epsilon \sim N(0, \sigma_{\omega}^2)$ with $\rho = 0.5$ and $\sigma_{\omega} = 1.05$. The first parameter is the persistence of hours documented by Mulligan (1997) while the second parameter replicates a Gini coefficient of earnings of around 0.6, as reported by Diaz-Gimenez, Glover and Rios-Rull (2011). The AR(1) process is then discretized using Tauchen's method to create $F(\omega'|\omega)$. Figure 2 shows the Lorenz curve of earnings and transfers, b , for this benchmark calibration.

A second version sets $\sigma = \eta = 1.5$ which is the standard assumption used, for example, by Castaneda, Diaz-Gimenez and Rios-Rull (2003). In order to keep the two calibrated versions comparable, we adjust α to 0.323 so that average transfers are the same in both models. Figure 2 also

shows the Lorenz curve of transfers for the alternative calibration. Clearly, the benchmark model, the one that differentiates between the EGS and the EIS, is able to produce more concentration and dispersion of wealth. The reason is that the standard model with low EGS and low EIS introduces too much aversion to consumption fluctuations, inducing too much savings, which tends to eliminate inequality. A higher EGS makes the prospect of falling into poverty less painful, or hitting the zero bequest constraint less problematic, and therefore a larger fraction of the population ends up hitting the constraint.

6 Concluding comments

The EIS has always played an important role in most macroeconomic models, determining both decisions within the lifetime of an individual, as well as across generations. This key role is in part due to the artifact that existing models assume the EIS and EGS to be identical. Once these concepts are disentangled, some of roles previously played by the EIS now belong to the EGS. For instance, we have shown how the EGS is a key determinant of the long-term fertility rate and long-run inequality. There are also instances in which the EGS is likely to play an important role in the short term. For instance, at the business cycles frequency, the EGS determines how a shock to the family budget, say an unemployment shock or winning the lottery, affects expenditures in children and, in particular, investments in their education and human capital formation.

Our paper is the first to formally model a distinction between intertemporal and intergenerational substitution. The utility representation we propose easily allows to associate a single parameter with the EIS, and a different one with the EGS. The simplicity of our preferences provides a useful and general framework for analyzing intergenerational issues. We expect this framework to introduce a new perspective, and to be useful in analyzing a number of interesting and relevant questions in economics.

References

- [1] Altonji, J., F. Hayashi and L. Kotlikoff. 1997. "Parental Altruism and Inter Vivos Transfers: Theory and Evidence." *Journal of Political Economy*, 105 (6), 1121-1166.
- [2] Barro, R. and G. Becker. 1989. "Fertility Choice in a Model of Economic Growth." *Econometrica*, 57 (2), 481-501.
- [3] Becker, G., and R. Barro. 1988. "A Reformulation of the Economic Theory of Fertility." *Quarterly Journal of Economics*, 103 (1), 1-25.
- [4] Becker, G., T. Philipson and R. Soares. 2005. "The Quantity and Quality of Life and the Evolution of World Inequality." *American Economic Review*, 95 (1), 277-291.
- [5] Biederman, D.K. and C.F. Goenner. 2008. "A life-cycle approach to the intertemporal elasticity of substitution." *Journal of Macroeconomics*, 30, 481-498.

- [6] Birchenall, J. and R. Soares. 2009. "Altruism, fertility and the welfare of children: Health policy evaluation and intergenerational welfare." *Journal of Public Economics*, 93, 280–295.
- [7] Castaneda, A., J. Diaz-Gimenez and V. Rios-Rull. 2003. "Accounting for the U.S. Earnings and Wealth Inequality." *Journal of Political Economy*, 111(4), 818-857.
- [8] Cordoba, J. C. and M. Ripoll. 2014. "Intergenerational Transfers and the Fertility-Income Relationship." Conditionally accepted, *Economic Journal*.
- [9] Cordoba, J. C. and M. Ripoll. 2011. "A Contribution to the Economic Theory of Fertility." Manuscript.
- [10] Diaz-Gimenez, J., A. Glover and V. Rios-Rull. 2011. "Facts on the Distributions of Earnings, Income, and Wealth in the United States: 2007 Update." *Federal Reserve Bank of Minneapolis, Quaterly Review*, February.
- [11] Dickie, M. and V. Messman. 2004. "Parental altruism and the value of avoiding acute illness: are kids worth more than parents?" *Journal of Environmental Economics and Management*, 48, 1146-1174.
- [12] Doepke, M. 2005. "Child Mortality and Fertility Decline: Does the Barro-Becker Model Fit the Facts?" *Journal of Population Economics*, 18 (2), 337-366.
- [13] Doepke, M. 2004. "Accounting for Fertility Decline During the Transition to Growth." *Journal of Economic Growth*, 9, 347-383.
- [14] Epstein, L., and S. E. Zin. 1989. "Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns I: A Theoretical Framework." *Econometrica*, 57 (4), 937-969.
- [15] Folbre, N. 2008. *Valuing Children: Rethinking the Economics of the Family*. Harvard University Press.
- [16] Golosov, M., L. Jones and M. Tertilt. 2007. "Efficiency with Endogenous Population Growth." *Econometrica*, 75 (4), 1039-1071.
- [17] Guryan, J., E. Hurst and M. Kearney. 2008. "Parental Education and Parental Time with Children." *Journal of Economic Perspectives*, 22 (3), 23-46.
- [18] Guvenen, F. 2006. "Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective." *Journal of Monetary Economics*, 53, 1451–1472.
- [19] Hall, R. and C. Jones. 2007. "The Value of Life and the Rise in Health Spending." *Quarterly Journal of Economics*, 122 (1): 39-72.
- [20] Hurd, M. and J. Smith. 2001. "Anticipated and Actual Bequest." In: Wise, D. (Ed.), *Themes in the Economics of Aging*, 357-392, National Bureau of Economic Research, University of Chicago Press.

- [21] Jones, C. and P. Klenow. 2011. "Beyond GDP? Welfare Across Countries and Time." NBER Working Paper No. 16352.
- [22] Jones, L., A. Schoonbroodt and M. Tertilt. 2011. "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?" In: Shoven, J. (Ed.), *Demography and the Economy*, 43-100, National Bureau of Economic Research, University of Chicago Press.
- [23] Jones, L. and A. Schoonbroodt. 2010. "Complements versus Substitutes and Trends in Fertility Choice in Dynastic Models." *International Economic Review*, 51 (3), 671-699.
- [24] Jones, L. and A. Schoonbroodt. 2009. "Complements versus Substitutes and Trends in Fertility Choice in Dynastic Models: Appendix." Manuscript.
- [25] Jones, L. and M. Tertilt. 2008. "An Economic History of Fertility in the United States: 1826-1960." In: Rupert, P. (Ed.), *Frontiers of Family Economics*, Volume 1, 165-230, Emerald Group Publishing Limited.
- [26] Mulligan, C. 1997. *Parental Priorities and Economic Inequality*. University of Chicago Press.
- [27] Murphy, K. and R. Topel. 2006. "The Value of Health and Longevity." *Journal of Political Economy*, 114 (5), 871-904.
- [28] Rosen, S. 1988. "The Value of Changes in Life Expectancy." *Journal of Risk and Uncertainty*, 1 (3), 285-304.
- [29] USDA. 2012. *Expenditures on Children by Families, 2011*. Miscellaneous Publication Number 1528-2011.
- [30] Viscusi, W. K. and J. Aldy. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World." *The Journal of Risk and Uncertainty*, 27 (1): 5-76.
- [31] Weil, P. 1990. "Nonexpected Utility in Macroeconomics." *Quarterly Journal of Economics*, 105 (1), 29-42.