LUISS

Department of Economics and Finance – Course in Economics and Business

Class of Applied Statistics and Econometrics

The effect of healthcare development on a country's fertility rate

Supervisor:
Andrea Pozzi

Candidate:
Renato Arnò
284461

Academic year: 2024/2025

Abstract

For most of history, human societies have been characterized by high rates of fertility and poor health conditions. This trend reversed substantially starting from the second half of the 20th century and hasn't stopped ever since. Is this a mere correlation, or does healthcare development play a role in a family's decision to have children? What kind of relationship binds these two variables, and what are the economic implications which come from studying this phenomenon? This work performs an empirical analysis using available real-world data and an econometric approach that produces results based on such data, which will be used to answer all the above questions. These results will later be tested on a series of aspects to verify their reliability and their robustness to potential biases.

TABLE OF CONTENTS

Introduction 4
The economic question 4
Why it is important 5
The methodology 6
Chapter 1: Building and presenting the model 9
1.1 Analysis of the variables taken separately 9
1.1.1 Per capita physicians 9
1.1.2 Infant mortality rate 11
1.1.3 Per capita GDP 12
1.2 The final model 14
1.2.1 Showcase of the results 14
1.2.2 Economic implications 15
Chapter 2: Robustness checks 17
2.1 Comparison with alternative versions 17
2.3 Endogeneity concerns 19
2.3.1 Infant mortality rate under age 5 19
2.3.2 Female labour force participation rate 20
2.4.3 HDI 23
2.4.4 Urbanization 25
Chapter 3: Considerations about the model 27
3.1 Strengths 27
3.2 Weaknesses 28
Conclusions 30
Bibliography 32

Introduction

The economic question

For most of history, human societies (both rural and urban) have been characterized by both poor health conditions and high rates of fertility. Having many children was not only a means of obtaining additional help in owned farms or businesses, but also a way of fighting high mortality and low average life expectancy which would otherwise bring to the demise of the whole family. Indeed, although it is true that in the past millennia it was common for women to give birth to five or more children, population growth was not as intense since poor healthcare combined with the lack of knowledge of germ theory (which brought people to ignore basic hygienic customs such as washing hands) caused the early death of many offsprings before they could reach reproductive age.

Up until the first half of the 20^{th} century the world fertility rate was relatively high compared to modern standards. Then, after a last peak reached in the 1960s, there was a sustained decline that continues to this day. This more recent trend is shown in the following graph:

Global fertility rate from 1960 to today

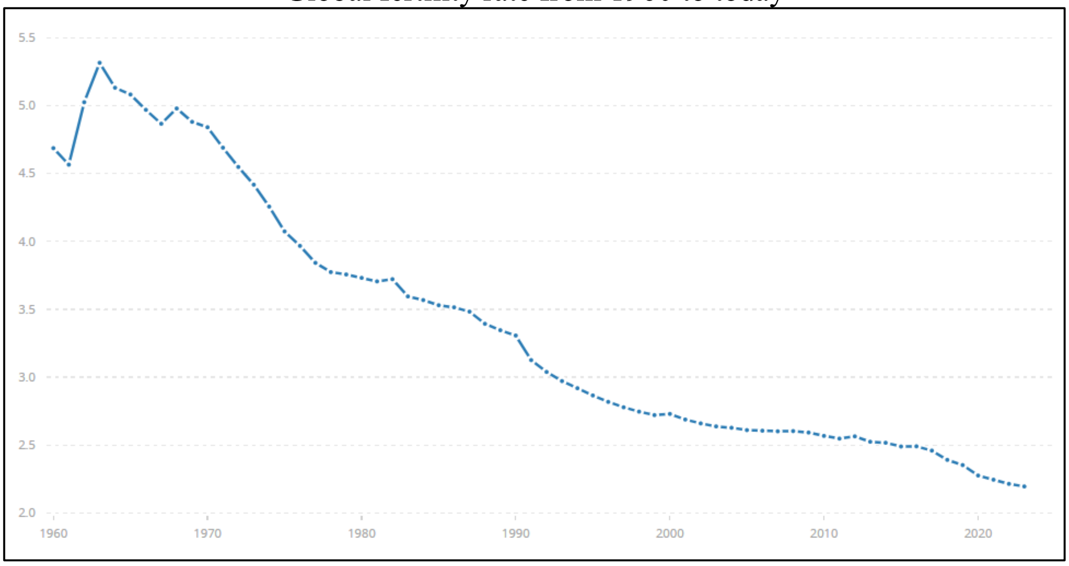

Source: World Bank Open Data

These years, however, were also accompanied by fast technological progress and economic development, which favoured an unprecedented improvement in medicine and in the overall standards of living, reducing mortality from various diseases and increasing average life expectancy. A large part of the current debate on what caused the recent decline in fertility advocates for the growth in income per capita, which fosters self-sufficiency and incentivizes the pursuit of individual success rather than setting up an enlarged family which would results in higher costs than benefits.

This work aims to move the attention on healthcare rather than income. In particular, it attempts to show that self-sufficiency driven by higher earnings plays a marginal role when the health conditions of potential offsprings are considered. The economic question is the following: what kind of role does development in healthcare have on a family's decision to have children? How relevant is its impact? Does current evidence suggest causality between these two factors?

The hypothesis to be tested is that as healthcare improves, the incentives to have more children decrease, since their probability of reaching adult age is higher. The objective of this research is *not* to demonstrate that this factor is the definitive cause of the recent phenomenon of declining births, nor to try and give any sort of complete answer to the problematic. The causes behind it are numerous and must not be disregarded. Rather, what this work attempts to do is establish a causal link between healthcare and fertility such that the former, together with other factors, acquires relevance in the current debate engaged in understanding the dynamics of the latter.

Why it is important

Studying the relationship between healthcare and fertility helps understand whether a country that improves its living conditions will inevitably face a shrinking population. Since development in health-related sectors is not only irreversible, but also likely to accelerate in the near future (as for anything linked to technological progress), it is important to acknowledge whether this will further worsen the current fertility situation, thus increasing the urgency for family-oriented policies. Alternatively, if governments are able to predict a decline in fertility in the coming years, they can more accurately adjust their policies to allocate resources in an efficient way in the long term on a series of sectors. To make some examples, a lower number of children will result in a lower demand for education, thus allowing for less spending on public schooling facilities and more elsewhere. Secondly, if a country cannot safely rely on a sustained growth of its own population, it could enact reforms favourable to immigration in an attempt to offset the reduction in births.

The methodology

To obtain relevant information for answering to the research question and verifying the initial hypothesis an empirical approach will be used, making use of tools and concepts coming from econometric theory. The first step is to gather data on both fertility and factors that could have a potential influence on it, with particular focus on healthcare. Subsequently, these data will be used to build a simple, cross-sectional OLS regression model with the aim of establishing a quantitative relationship between the two sides of the equation. This process requires two phases: first, each of the considered independent variables will be analysed individually with respect to fertility rate, so as to find the best type of function that binds them together (e.g. linear, logarithmic, inverse). Only after this initial step it will be possible to obtain an optimal multivariate model. For example, a variable which influences the outcome of interest exponentially will perform poorly if considered in linear form and will thus yield biased results.

Keeping this in mind, the best-performing model obtained at the end of the research efforts and that will be implemented to answer the main economic question has the following equation:

$$FertRate_i = \beta_0 + \beta_1 PcapPhys(Inv)_i + \beta_2 InfMortRate_i +$$

$$+\beta_3 PcapGDP(Inv)_i + \beta_4 Islamic_i$$

Where:

- $\boldsymbol{FertRate_i}$: Fertility rate (number of children per woman) in country *i;*

- $\boldsymbol{PcapPhys(Inv)_i}$: Inverse of per capita physicians (number of physicians per 1000 people) in country i. The term "physicians" refers to medical personnel, both specialized and non-specialized;

- $\boldsymbol{InfMortRate_i}$: Infant mortality rate under age 1 (number of newborns per 1000 that die before reaching the first year of age) in country *i*;

- $\boldsymbol{PcapGDP(Inv)_i}$: Inverse of per capita GDP (expressed in US dollars) in country *i*;

- $\boldsymbol{Islamic_i}$: A dummy that takes value 1 if the majority/official state religion in country i is islam and 0 otherwise;

The main explanatory variable of the model is the number of per capita physicians, which is used as a proxy for the level of healthcare development. A higher value of this variable should indicate a more developed healthcare sector. The infant mortality rate is instead a control variable that might have a potential correlation with per capita physicians. Indeed, a higher mortality of newborns can be a signal of lower healthcare development. But mortality can be also due to many other factors, such as wars, poverty and famines. Including this variable in the model will thus help substantially in reducing potential OVB. The same can be said for per capita GDP, the second control. A richer population is able to afford better healthcare facilities. This would mean that the effect being captured would not be that of a more developed healthcare sector, but rather one of a more advanced economy. It is important that the variation captured by this variable be isolated from the main regressor to further reduce biases in the estimates. The hypothesis behind the dummy that controls for Islam-majority societies is that in such places customs allow for early and arranged marriages, and women are typically enforced the role of taking care of the house and the family. These factors would contribute to higher-than-normal values of fertility rates and would thus constitute a bias in the measures. Hence, they must be checked for as well.

As can be noted, the analysis is performed at country level. The data for each variable have been collected from the large data repositories available on the World Bank Open Data website. The reference year is 2020 for each component in order to give a homogeneous time frame for the estimates. The countries observed are 100: 25 from Europe, Africa and Asia each; 23 from America, and 2 from Oceania. The sample was organized in this way to spread the observations across the world and get measures as heterogeneous and equally distributed as possible.

After presenting the model's estimates and making all the relevant comments in both statistical and economic fields, a series of tests to the robustness of these estimates will be carried out. These will include comparing the values of the original regression with the ones of alternative versions implementing combinations of different transformations in the considered variables, as well as addressing various endogeneity concerns by inserting new predictors correlated with already included ones, to see if there are any biases in the estimates. Finally, some further considerations will be made about the strengths and weaknesses of the model.

To obtain all the needed results, which include both numerical values and graphical representations on dedicated scatterplot graphs, the tool used is Python with its specialized libraries.

Chapter 1: Building and presenting the model

This chapter is dedicated to the construction and the showcase of the estimates of the regression model used to answer the economic question posed in the introduction.

1.1 Analysis of the variables taken separately

The study of the individual components of the model will be initially carried out. This is a necessary procedure to understand the kind of relationship that binds each observed factor to the dependent variable. Indeed, as will be shown, the choice of a linear function might not always be the most optimal. By the end of this section, it will be clear why each independent variable is plugged in the model equation in its own way.

1.1.1 Per capita physicians. The first guess for each individual regressor on fertility rate will be the linear relationship. In this case, Fig. 1.1 shows the output of an OLS regression considering only the number of per capita physicians, whereas Fig. 1.2 is the scatterplot graph derived from the data.

Figure 1.1 - Linear regression output

Dependent variable is fertility rate		Coefficient	Standard error	Z-statistic	p-value
Per capita physicians		-0.4842	0.050	-9.603	0.000
No. Observations	R^2	Adj. R^2	F-statistic	AIC	BIC
100	0.495	0.490	92.21	240.0	245.2

Figure 1.2 - Graphical representation

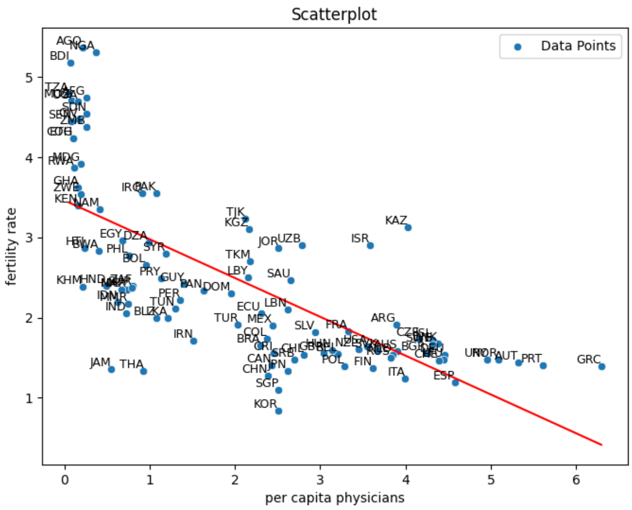

From the estimates it can be seen that the coefficient is highly significant. According to the results, a unit increase in per capita physicians is associated, on average, with an approximately 0.4842 decrease in fertility rate. The R-squared is 0.495, meaning that almost half of the variation in the dependent variable is explained by this model.

These results might seem satisfying. However, by looking at the scatterplot graph in Fig.1.2, the observations don't quite seem to be following a linear pattern. Indeed, the trend seems to be that of a curve. Because of this, the question arises for whether an alternative model may offer a better fit than the linear one. Since the shape resembles a hyperbole, an inversely proportional relationship could be a good guess. By taking the inverse of the values of per capita physicians and running a new regression, new estimates are obtained (Fig. 1.3). Furthermore, it is possible to draw a new scatterplot with an inversely proportional fit (Fig. 1.4).

Figure 1.3 - Inversely proportional regression output

Dependent variable is fertility rate		Coefficient	Standard error	Z-statistic	p-value
Inverse of per capita physicians		0.2423	0.035	6.995	0.000
No. Observations	R^2	Adj. R^2	F-statistic	AIC	BIC
100	0.542	0.538	48.93	230.1	235.3

Figure 1.4 - Graphical representation

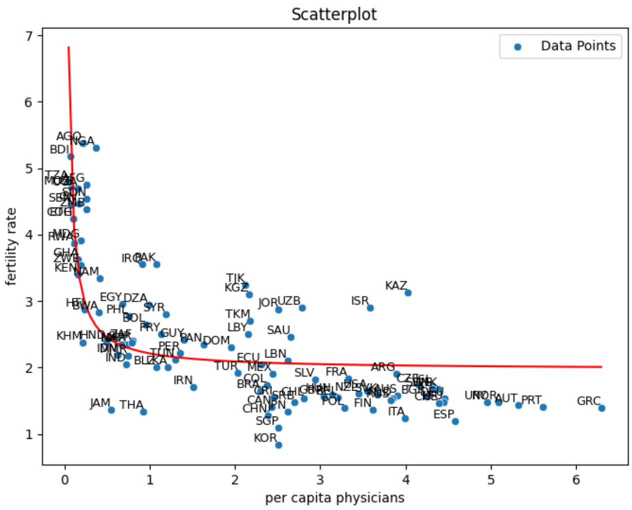

The results seem to confirm that inverse proportionality is a better suited explanator of the relationship between the two variables, as shown by the higher R-squared (0.542) and the lower AIC and BIC statistics (230.1 and 235.3 compared to 240 and 245.2 of the linear version). Moreover, from the graph in Fig. 1.4 it can be seen that the hyperbole better follows the trend of the observations. Indeed, it seems that as the number of per capita physicians increases, the effect of one additional physician on fertility rate becomes less relevant. This is expressed by the progressive flattening of the curve. According to the results, a unit increase in the inverse of per capita physicians is associated, on average, with an approximately 0.2423 increase in fertility rate. This is why in the equation of the model shown in the introduction this variable is expressed in its inverse transformation.

1.1.2 Infant mortality rate. The impact of infant mortality rate is now examined. Fig. 1.5 shows the output of an OLS regression considering only this variable, whereas Fig. 1.6 is the scatterplot graph derived from the data.

Figure 1.5 - Linear regression output

Dependent variable is fertility rate		Coefficient	Standard error	Z-statistic	p-value
Infant mortality rate		0.0622	0.004	14.106	0.000
No. Observations	R^2	Adj. R^2	F-statistic	AIC	BIC
100	0.748	0.746	199.0	170.3	175.5

Figure 1.6 - Graphical representation

From the estimates it can be seen that the coefficient is highly significant. According to the results, a unit increase in infant mortality rate is associated, on average, with an approximately 0.0622 increase in fertility rate. The R-squared is 0.748, meaning that almost three quarters of the variation in the dependent variable are explained by this model.

The trend of the observations in the scatterplot seems to follow a straight line. Indeed, in this case the linear relationship provides a good fit for the data. This is why in the equation of the model shown in the introduction, the variable of infant mortality rate is expressed in the standard linear form.

1.1.3 Per capita GDP. The next variable analysed is per capita GDP (expressed in US dollars). Following the procedure applied thus far, the linear fit is attempted as a first guess. To work with smaller numbers (and thus avoid long sequences zeros in the estimates), the values of this variable have been divided by 1000. Figures 1.7 and 1.8 show the regression output and the graphical visualization of the data, respectively.

Figure 1.7 - Linear regression output

Dependent variable is fertility rate		Coefficient	Standard error	Z-statistic	p-value
Per capita GDP		-0.0304	0.004	-7.505	0.000
No. Observations	R^2	Adj. R^2	F-statistic	AIC	BIC
100	0.293	0.286	56.33	273.5	278.8

Figure 1.8 - Graphical representation

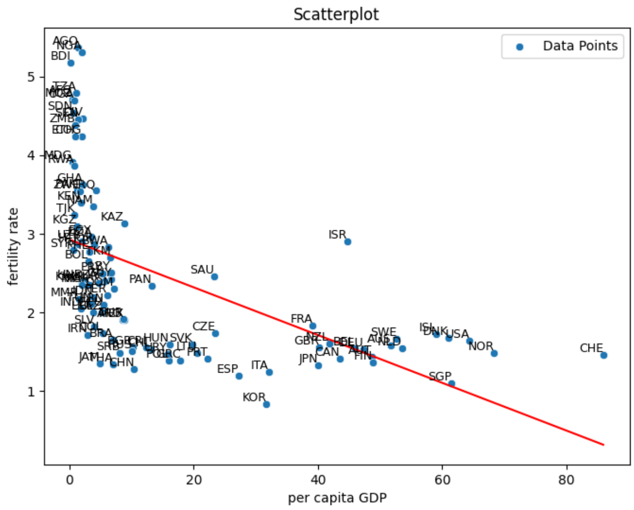

From the estimates it can be seen that the coefficient is highly significant. According to the results, a unit increase in thousands of per capita GDP is associated, on average, with an approximately 0.0304 decrease in fertility rate. The R-squared is 0.293, meaning that 29.3% of the variation in the dependent variable is explained by this model. This low value might be due to linearity not being an optimal fit for the data. This is confirmed by what is seen on the scatterplot graph. As in the case of per capita physicians, the disposition of the observations seems to follow a hyperbolic shape. Thus, an inversely proportional relationship could be a better guess. By taking the inverse of the values and running a new regression, new estimates are obtained (Fig. 1.9). Furthermore, it is possible to draw a new scatterplot with an inversely proportional fit (Fig. 1.10).

Figure 1.9 - Inversely proportional regression output

Dependent variable is fertility rate		Coefficient	Standard error	Z-statistic	p-value
Inverse of per capita GDP		1.1721	0.263	4.452	0.000
No. Observations	R^2	Adj. R^2	F-statistic	AIC	BIC
100	0.467	0.462	19.82	245.2	250.5

Figure 1.10 - Graphical representation

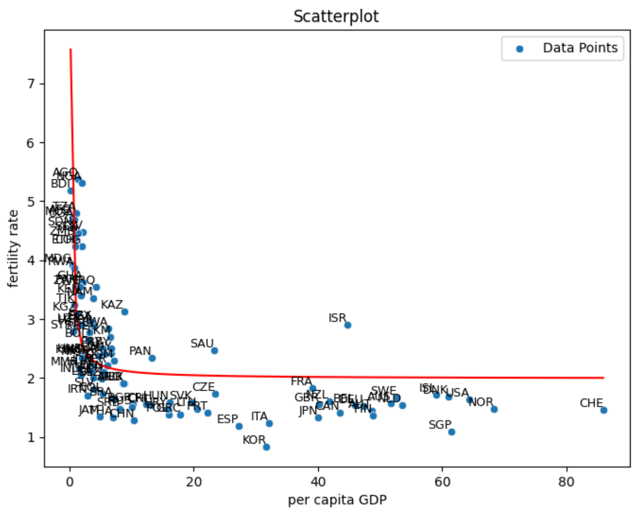

The results seem to confirm that inverse proportionality is a better suited explanator of the relationship between the two variables, as shown by the higher R-squared (0.467) and the lower AIC and BIC statistics (245.2 and 250.5 compared to the 273.5 and 278.8 of the linear version). Moreover, from the graph in Fig. 2.14 it can be seen that the hyperbole better follows the trend of the observations. Indeed, it seems that as per capita GDP increases, the effect on fertility rate becomes less relevant. This is expressed by the progressive flattening of the curve. According to the results, a unit increase in the inverse of thousands of per capita GDP is associated, on average, with an approximately 1.1721 increase in fertility rate. This is why in the equation of the model shown in the introduction, this variable is expressed in its inverse transformation.

1.2 The final model

Now that all the factors have been studied singularly and their relationship with the dependent variable figured out, they are ready to be implemented in the multivariate regression. This section shows the results of the best-performing model (which corresponds to the one shown in the introduction) obtained at the end of the research efforts.

1.2.1 Showcase of the results. The output of the regression which offers the best fit for the data is shown in Fig. 1.11.

Figure 1.11 - Estimates of the full model

Dependent variable is fertility rate		Coefficient	Standard error	Z-statistic	p-value
Inverse of per capita physicians		0.1115	0.014	8.063	0.000
Infant mortality rate		0.0411	0.005	7.797	0.000
Inverse of per capita GDP		0.1670	0.060	2.767	0.006
Islam dummy		0.4551	0.117	3.890	0.000
No. Observations	R^2	Adj. R^2	F-statistic	AIC	BIC
100	0.858	0.852	217.0	118.9	131.9

First, it must be noted that all the coefficients are highly significant, with all p-values strictly smaller than 0.01. In particular, the main explanatory variable (inverse of per capita physicians) manages to retain high significance even after controlling for multiple correlated factors.

According to the results, a unit increase in the inverse of per capita physicians (keeping all other factors unchanged) is associated, on average, with an approximately 0.1115 increase in fertility rate. The R-squared is 0.858, meaning that more than 85% of the variation in the dependent variable is explained by this model. The F-statistic is above 200, giving further proof of the significance of the implemented variables and their ability to offer an optimal fit for the data.

1.2.2 Economic implications. These results provide relevant information for answering the economic question posed in the introduction. In particular, if the number of per capita physicians is an accurate proxy for healthcare development, the interpretation of the estimates obtained is that this factor has an undeniable influence on the fertility rate of a country. However, by the nature of the inversely proportional relationship which binds the two variables, the magnitude of this effect reduces progressively as healthcare improves. This indicates that a better healthcare system influences the decision of women to have children less if a certain level of development is already achieved, compared to when conditions are more critical. The effect loses virtually all its relevance at fertility rates between 1 and 2 children per woman. This means that if a country were to invest its resources to achieve higher healthcare standards, it would inevitably face a slow decline in its population in the long term, since the minimum fertility rate necessary for a stable demographic growth is 2.1 children per woman.

Currently, there is no evidence that a declining/aging population has a direct negative effect on a country's economic growth[1]. However, the concern rises in the very long term, where the scarcity of young labour force combined with a high number of old individuals will put a strain on state finances through reduced tax revenues and increased pension expenditures. This doesn't mean that a purposeful deterioration of the healthcare system is the solution. As said, this factor becomes almost irrelevant after a certain point. Instead, it is necessary to increase the level of commitment into adopting a series of family-oriented policies to incentivize motherhood, while at the same time maintaining constant efforts in improving health conditions inside the society. This will make it easier to set up larger families but will also guarantee a brighter future for the new generations by minimizing the probability of growing unhealthy and dying at an early age.

[1] Daron Acemoglu, Pascual Restrepo: "Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation", AER, 2017

Chapter 2: Robustness checks

This chapter focuses on challenging the robustness of the model by comparing its estimates with the ones of similar versions using the implemented variables expressed in a different form. After this, additional factors will be considered which are suspected of sharing a potential correlation with one (or more) predictors. If the suspects are confirmed, these variables should undermine the significance of such predictors and thus identify a fallacy in the model.

2.1 Comparison with alternative versions

The model presented in chapter 1 offers the best results of all the ones tested. To show this, Fig. 2.1 provides the outcomes of alternative versions that use the same independent variables in their logarithmic transformations, whereas in Fig. 2.2 such change is applied to fertility rate instead.

Figure 2.1 - Transformations in the independent variables

Dependent variable is fertility rate	R^2	Adj. R^2	F-statistic	AIC	BIC
All log	0.764	0.754	91.39	170.0	183.0
Log physicians and infant mortality rate	0.790	0.781	138.2	158.4	171.4
Log physicians and GDP	0.821	0.813	115.0	142.5	155.5
Log infant mortality rate and GDP	0.797	0.789	91.36	154.6	167.6
Log physicians	0.836	0.829	201.3	133.4	146.4
Log infant mortality rate	0.809	0.801	127.8	148.7	161.7
Log GDP	0.855	0.849	160.4	121.4	134.4
Original model	0.858	0.852	217.0	118.9	131.9

Figure 2.2 - Transformation of fertility rate

Dependent variable is log fertility rate	R^2	Adj. R^2	F-statistic	AIC	BIC
All log	0.788	0.779	127.9	-32.85	-19.83
Log physicians and infant mortality rate	0.800	0.791	172.9	-38.75	-25.72
Log physicians and GDP	0.809	0.801	142.6	-43.69	-30.66
Log infant mortality rate and GDP	0.808	0.800	113.4	-43.15	-30.12
Log physicians	0.812	0.804	185.8	-45.11	-32.09
Log infant mortality rate	0.813	0.805	140.9	-45.54	-32.51
Log GDP	0.821	0.814	158.3	-50.14	-37.11
Without logs	0.814	0.806	143.6	-46.24	-33.22

The estimates in Fig. 2.1 do in fact confirm the superiority of the original model over its logarithmic variations, which present in all cases lower values of R-squared, adjusted R-squared and F-statistics. The same conclusion can be reached by looking at the AIC and BIC values, useful for comparing models based on the same dataset. In fact, the version that doesn't contain any logarithmic variables presents both the lowest AIC (118.9) and BIC (131.9), meaning that this is the build which offers the best trade-off between complexity and prediction capacity. One interesting aspect encountered during this analysis is that when isolated, the number of per capita physicians actually better explains the variation in fertility rate in a logarithmic relationship instead of an inversely proportional one. However, as shown in Fig. 2.1, the multivariate model with this change is still outperformed by the original one.

Moving on to Fig. 2.2 it can be noted that the best performing model is the one where only per capita GDP is transformed. However, it seems like the predictors implemented are better suited to explain absolute variations in the dependent variable, rather than percentage changes. Indeed, the best fit obtained by regressing on log fertility rate has an R-squared and F-statistic of respectively 0.821 and 158.3, way lower than the 0.858 and 217.0 of the original model. AIC and BIC statistics cannot be used to compare models predicting different dependent variables.

2.3 Endogeneity concerns

2.3.1 Infant mortality rate under age 5. A first reasonable endogeneity concern can arise over the variable of infant mortality rate used. The one in the original model takes into account only the deaths of children under the first year of age. A mortality rate extended up to age 5 could perhaps prove to be a more accurate control, since the one used so far is a subset of the latter (children dying before age 1 automatically contribute to the age 5 mortality rate), thus potentially omitting a large amount of information. Fig. 2.3 gives a quick look on the most fit relationship of this variable with fertility rate (in this case, linear), whereas Fig. 2.4 shows the values obtained by including it in the model in both linear and logarithmic forms. Instead of the coefficient estimate, the p-value of each predictor is reported since the focus of this analysis is to assess the significance of the independent variables, rather than the magnitude of their effect.

Figure 2.3 - Graphical representation

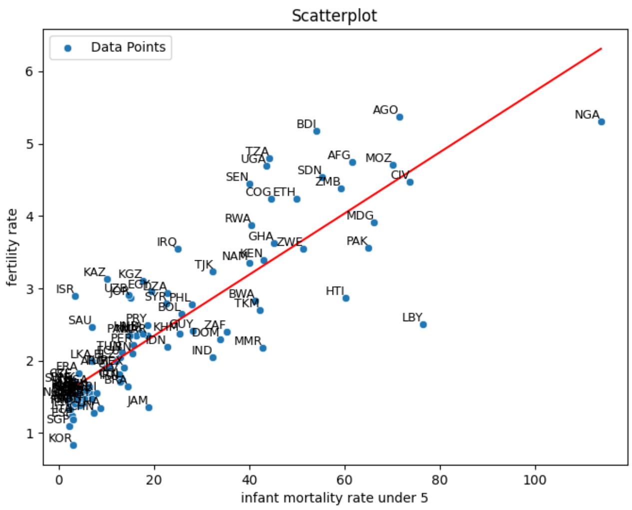

Figure 2.4 - p-values

Dependent variable is fertility rate	OLS (1)	OLS (2)
Inverse of per capita physicians	0.000	0.000
Infant mortality rate (under 1)	0.016	0.000
Infant mortality rate (under 5)	0.116	-
Log Infant mortality rate (under 5)	-	0.489
Inverse of per capita GDP	0.004	0.006
Islam dummy	0.000	0.000

If the endogeneity concern turned out to be true, the significance of infant mortality rate under age 1 should have dropped drastically. Instead, it was only reduced to just below the 1% level (furthermore, there were no meaningful changes in the other factors). Indeed, the only variable that seems to be out of place is the very one that was just added, not even being significant at the 10% level. A possible counterargument to this conclusion could be that the best fit for this variable taken separately might not be the same in a multivariate model, as was the case with per capita physicians. However, as shown, by attempting to include the logarithmic transformation of infant mortality rate under 5 in the regression the same results are obtained. Also, its p-value becomes even higher (0.489 compared to the 0.116 in linear form). This hypothesis of endogeneity can thus be confidently discarded.

2.3.2 Female labour force participation rate. The model that was considered best performing before the final one did not contain per capita GDP nor the Islam dummy. In their place was female labour participation rate, evaluated as the percentage of female population in a country aged 15+ participating to the labour force. When analysed separately, the mathematical function best expressing the relationship with fertility rate had the following structure:

$$FertRate_i = \beta_0 + \beta_1 FPartRate_i^3 + \beta_2 FPartRate_i$$

With β_1>0 and β_2<0. The graphical representation of this relationship in a scatterplot is shown in Fig. 2.5

Figure 2.5 - Graphical rapresentation

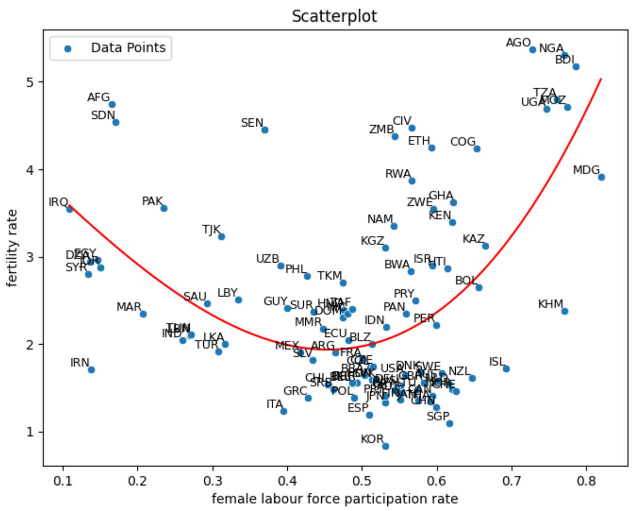

To both low and high values of participation rate correspond high fertility rates, whereas the lower ones are concentrated around values between 0.5 and 0.6 of the independent variable. Fig. 2.6 reports the output of the multivariate regression including this predictor.

Figure 2.6 - Regression output

Dependent variable is fertility rate	Coefficient	Standard error	Z-statistic	p-value
Inverse of per capita physicians	0.1110	0.018	6.287	0.000
Infant mortality rate	0.0434	0.004	9.971	0.000
Female labour force participation rate (cubic component)	3.2196	1.249	2.578	0.010
Female labour force participation rate (linear component)	-2.6943	0.787	-3.421	0.001

No. Observations	R^2	Adj. R^2	F-statistic	AIC	BIC
100	0.844	0.838	125.4	128.3	141.3

This model proved to withstand the robustness-checking phase (also managing to undermine the significance of per capita GDP) up until the Islam dummy was added. Indeed, as it can be noted in Fig. 2.5, the observations concentrated on low values of participation rate correspond to countries where Islam is the majority/state official religion (Iran, Afghanistan, Syria, Lybia, Saudi Arabia etc). Hence, it comes as no surprise that, when controlling for this factor, the original polynomial relationship "breaks". However, it has no major impact on per capita GDP, thus reinstating its significance in the new, improved model which would be designated as the definitive one. Fig. 2.7 shows the variables' p-values in multiple regressions including different combinations of the female labour force participation rate components inside the final model.

Figure 2.7 - p-values

Dependent variable is fertility rate	OLS (1)	OLS (2)	OLS (3)
Inverse of per capita physicians	0.000	0.000	0.000
Infant mortality rate	0.000	0.000	0.000
Female labour force participation rate (cubic component)	0.247	0.153	-
Female labour force participation rate (linear component)	0.530	-	0.311
Inverse of per capita GDP	0.129	0.038	0.008
Islam dummy	0.003	0.000	0.001

If the evaluation criteria were simply that if the p-value rises above 0.1 the variable should be dropped, these results would be of ambiguous interpretation. Even though it is true that, by adding female labour force participation rate, per capita GDP loses significance at the 10% level, it must also be noted that the inserted components themselves are not significant (by an even larger margin). The two factors are competing in explaining the same portion of the variation in fertility rate, and this can be due to multicollinearity, which inflates standard errors and renders the estimates less reliable. A logical interpretation of the results reported in Fig. 2.7 is that per capita GDP is not statistically insignificant *per se*, but rather its standard error is altered by the introduction of female labour force participation rate in the model, which *is* statistically insignificant *per se*, an outcome which wouldn't have occurred had the Islam dummy (which is *not* collinear with per capita GDP) not been inserted. This would also explain why the p-value of per capita GDP is still smaller than the one of the two components of the added variable. To test this, checking the VIFs (Variance Inflation Factors) of the predictors would turn out useful. This statistic helps establish the level of multicollinearity of each variable. Generally, concerns should be raised over values above 5. Fig. 2.8 shows a table with all the VIFs calculated on different models.

Figure 2.8 - Variance Inflation Factors

Dependent variable is fertility rate	VIF (1)	VIF (2)	VIF (3)	VIF (4)	VIF (5)	VIF (6)
Inverse of per capita physicians	2.287083	-	-	-	1.482037	-
Infant mortality rate	1.882410	-	1.312145	-	-	-
Female labour force participation rate (cubic component)	10.698329	6.074408	7.895534	9.497951	7.954073	7.204087
Female labour force participation rate (linear component)	12.109386	6.074408	7.334026	8.531738	6.776045	9.909848
Inverse of per capita GDP	2.313362	-	-	1.578865	-	-
Islam dummy	2.238871	-	-	-	-	2.160109

The results confirm the initial hypothesis. Indeed, the values for both the cubic and linear terms of female labour force participation rate are higher than 10 (thus indicating strong multicollinearity), whereas the one for per capita GDP is only slightly greater than 2. It is also worth noting that when considered without other predictors, the polynomial components already show VIFs above 5, but both increase substantially when per capita GDP (and nothing else) is added (VIF(4)). This is the definitive proof that yes, the two components are first of all collinear with each other (which could be expected, since one is simply the cubic transformation of the other), but collinearity also runs towards per capita GDP, and *not* the other way around, making female labour force participation rate not a reliable tool to undermine the significance of the already included variable. This is why the original model has once again proven its robustness even after an apparent initial fallacy.

2.4.3 HDI. Another possible endogeneity concern might arise by not accounting for a country's Human Development Index (HDI). The hypothesis is that other than a potential correlation with per capita GDP, places with higher HDIs might dispose of a more developed healthcare sector. The best fit on fertility rate is given by an inversely proportional relationship (Fig. 2.9) even though not by a large margin from the linear and logarithmic one. This is why all three versions are tried when testing the robustness of the original model. The p-values of the coefficients in all the tested regressions are shown in Fig. 2.10.

Figure 2.9 - Graphical representation

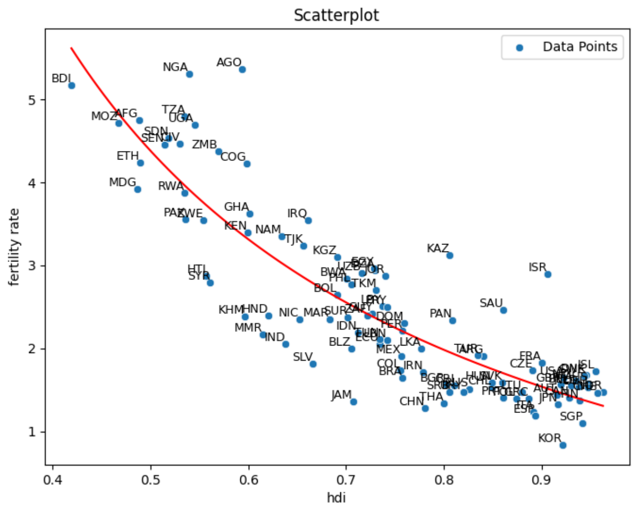

Figure 2.10 - p.values

Dependent variable is fertility rate	OLS (1)	OLS (2)	OLS (3)
Inverse of per capita physicians	0.000	0.000	0.000
Infant mortality rate	0.000	0.000	0.000
HDI	0.897	-	-
Inverse of HDI	-	0.983	-
Log HDI	-	-	0.911
Inverse of per capita GDP	0.004	0.025	0.008
Islam dummy	0.000	0.000	0.000

Even though from Fig. 2.10 this variable seems to be a more than valid explanator of the variation in fertility rate (the R-squared of the corresponding regression was indeed 0.760), its significance plummets when inserted in the multivariate model, both in logarithmic and linear form, leaving the p-values of all already present predictors almost unchanged. This hypothesis of endogeneity can thus be confidently discarded.

2.4.4 Urbanization. The next robustness check is performed by adding a control for the urbanization rate in the observed country, expressed as the percentage of total population that lives in urban areas. The endogeneity concern arises over the fact that lower values indicate more rural societies in which workers are self-employed by working their own land. This might have potential effects on GDP (lack of high-skill labour translates into lower productivity and thus lower income) and the supply of physicians. But also, rural families should have more incentives to have children due to their flexible working schedule (which does not impose time constraints on the parents) and the need for additional help in manual labour. Fig. 2.11 shows the scatterplot graph representing the best fit for this variable analysed separately, whereas Fig.2.12 reports the p-values of the usual tested multivariate models.

Figure 2.11 - Graphical representation

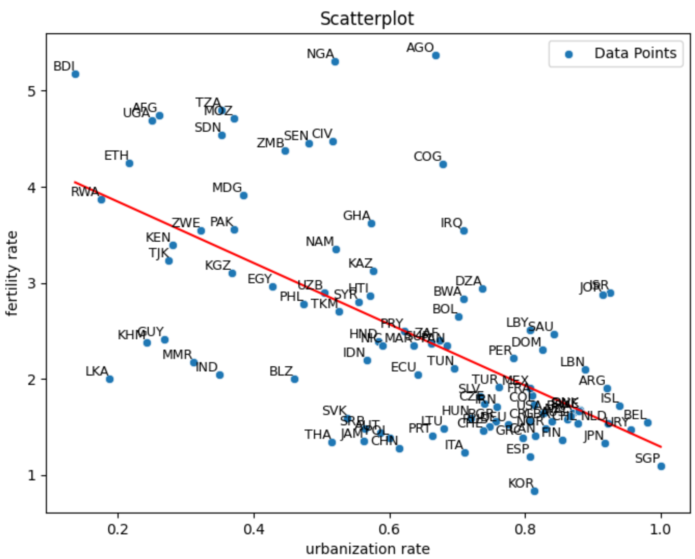

Figure 2.12 - p-values

Dependent variable is fertility rate	OLS (1)	OLS (2)	OLS (3)
Inverse of per capita physicians	0.000	0.000	0.000
Infant mortality rate	0.000	0.000	0.000
Urbanization rate	0.652	-	-
Inverse of Urbanization rate	-	0.961	-
Log Urbanization rate	-	-	0.791
Inverse of per capita GDP	0.003	0.020	0.005
Islam dummy	0.000	0.000	0.000

Once again, all the coefficients except the one of the newly included variable retain low p-values. This indicates that even when controlling for urbanization, the other predictors don't lose their significance, proving the overall robustness of the original model. Thus, this endogeneity concern can be confidently discarded as well.

Chapter 3: Considerations about the model

Having conducted a thorough analysis of the results provided by the implemented regression model in terms of accuracy, economic implications and robustness, this short final chapter now makes a series of impartial considerations on the effective strengths and weaknesses of such model.

3.1 Strengths

An initial assessment of the strengths can be carried out by looking at the statistics obtained by running the main regression. First it must be noted again that all the coefficients are significant over the 1% level by a large margin, especially the one of the main explanatory variable. The model has successfully answered the economic question posed in the introduction, leaving no uncertainty about the relevance of healthcare in influencing fertility. The R-squared of 0.858 conveys that the implemented predictors are able to explain a large chunk of the variation in the independent variable. Moreover, the adjusted version being close to the original (0.852) means that the number of independent variables included does not add unnecessary complexity in predicting the values of fertility rate. This specific model even provides the best trade-off between complexity and prediction-accuracy of any other alternative version tested, with its AIC and BIC being the lowest among all its counterparts (except the ones that predict log fertility rate, but as already mentioned in section 2.1, model comparison with AIC and BIC values is not possible with different dependent variables. For these cases, their inferior performance is determined by their lower R-squared and F-statistics). Multicollinearity does not seem to be a problem (since standard errors are not excessively high relative to the estimates), and the VIFs of each variable shown in Fig. 3.1 further confirms this.

Figure 3.1 - Variance Inflation Factors

Variable	VIF
Inverse of per capita physicians	2.070324
Infant mortality rate	1.312145
Inverse of per capita GDP	1.996862
Islam dummy	1.099512

Indeed, all the values are strictly below 2.1, indicating extremely low multicollinearity.

As shown in chapter 2, the model has also proven its robustness to multiple endogeneity threats. Every time the insertion of a new variable was attempted, the portion of variation in fertility rate it was initially explaining alone would be captured by the other predictors, making it lose its significance (and never the other way around).

3.2 Weaknesses

Even though the model presents many strengths, from its statistical accuracy to its resistance to endogeneity tests, there are also a few meaningful weaknesses to be addressed.

First, the number of observations is not among the highest. This is not to say that the estimates obtained are unreliable, but rather that they lack some precision due to a restricted reference sample. The reason behind analysing only 100 countries despite the existence of many more in the real world is the difficulty of retrieving complete and homogeneous data on multiple variables for each, especially the ones currently undergoing internal crises such as conflicts or oppressive governments.

A second weakness is the simplicity of the model. A standard cross-sectional OLS regression fails to capture changes in the variables over prolonged periods of time, since it relies on stationary measures (in this case, only the year 2020 was considered). The main challenge remains the collection of exhaustive data for each observation.

Other than in the model structure, weaknesses might be found also in the choice of the main predictor. Precisely, there is the possibility that the number of per capita physicians is not an ideal proxy for the level of healthcare development of a country, since it just measures the amount of medical personnel, and not the quality of such personnel or of healthcare facilities. It is true, however, that it has demonstrated its superiority over other possible proxies considered, such as per capita health expenditure in terms of GDP and the number of hospital beds per 1000 people. Nevertheless, the reliability of these alternative factors is itself unknown, making the interpretation ambiguous: either they are reliable (and per capita physicians is not), thus undermining the effect of healthcare development on fertility rate (since they are not that significant); or, they are not reliable, and per capita physicians is the best proxy available for this factor. The only certain thing is that its effect on fertility rate is highly significant and unaffected by other correlated variables. Whether this effect corresponds to healthcare development cannot be established confidently. However, this remains the most logical interpretation to adopt.

Conclusions

Does healthcare development have an effect on fertility? The analysis seems to suggest that the answer is yes. Is this effect linear? From a quantitative perspective, the data show high significance of this relationship. However, better results are obtained when attempting an inversely proportional fit. The methodology implemented in this work argues that for this reason, it can be concluded that the latter is better suited for binding the two variables. This claim is in fact reinforced when all other predictors are considered. Naturally, this does not mean that a variable will necessarily behave the same way both isolated and in a multivariate regression. For example, when the logarithmic transformation of per capita physicians was attempted, it initially had the highest explanatory power when analysed separately, but it ultimately proved to be inferior to its inversely proportional counterpart in the complete model. It is true, however, that all these alternative ways of expressing the main variable give the same outcome of high significance, only in different intensities. This suggests that the effect of healthcare on fertility is present and is powerful, so much so that it can withstand inaccuracies in the way it is measured.

With the coefficients being highly significant at the 1% level, the R-squared over 0.85 and endogeneity highly unlikely, the estimates of the regression model used provide strong and unbiased evidence in the support of the initial hypothesis. Furthermore, they add more detail to the dynamics of the relationship between the two factors of interest by showing that as healthcare improves, its impact on fertility decreases in intensity.

In economic terms, the results convey valuable information on possible long-term oriented policies that governments should undertake. If their objectives are to adapt to the future trends in fertility, resources shall be allocated more on sectors that are reliant on population growth, such as the available stock of workforce and pensions, and less on the ones that are adversely affected by it, such as education (fewer children reduce the need for additional schooling facilities). Alternatively, knowing that healthcare loses relevance over fertility values between 1 and 2, governments can increase their commitments to incentivise more numerous families without worrying about further negative effects caused by improved health conditions.

The answer to the research question drawn at the end of the analysis holds on a series of assumptions on the methodology used, such as the sufficiency of the tests performed on both the transformations of the variables and their robustness to endogeneity, the accuracy of per capita physicians as a proxy for healthcare development, and the reliability of a cross-sectional, time-invariant regression estimates. This work enforces its conclusions only to the extent that these postulates are assumed to be true and does not claim that the *modus operandi* adopted is flawless. It just shows that, given the dataset and the methodology used, these are the results obtained.

Bibliography

https://data.worldbank.org/indicator/SP.DYN.TFRT.IN

https://data.worldbank.org/indicator/SH.MED.PHYS.ZS

https://data.worldbank.org/indicator/SP.DYN.IMRT.IN

https://data.worldbank.org/indicator/NY.GDP.PCAP.CD

https://data.worldbank.org/indicator/SH.DYN.MORT

https://data.worldbank.org/indicator/SL.TLF.CACT.FE.ZS

https://countryeconomy.com/hdi?year=2020

https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS

https://data.worldbank.org/indicator/SH.XPD.CHEX.PC.CD

https://data.worldbank.org/indicator/SH.MED.BEDS.ZS

Daron Acemoglu, Pascual Restrepo: “Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation”, AER, 2017