2014

A Bayesian Model of Fertility Decisions in Relationship to Female Labor Force Participation

Rebecca C. Wardrop
The College of Wooster, rcwardrop@gmail.com

A Bayesian Model of Fertility Decisions in Relationship to Female Labor Force Participation

By
Rebecca Wardrop

Submitted in Partial Fulfillment of the Requirements of Senior Independent Study for the Department of Economics and the Department of Mathematics at the College of Wooster

Advised by
Dr. Barbara Burnell
Department of Economics and
Dr. James Hartman
Department of Mathematics

24 March 2014

Abstract

Due to the increasing number of women in the labor force, opportunity costs associated with labor force participation are becoming an important factor in fertility decisions. Further, these decisions are assumed to be dynamic as the opportunity costs change as a woman progresses through her career. A Bayesian statistical model , which allows the distribution of the likelihood of having children to be updated as information is gathered, lends itself to the dynamicity of the decision-making process. A generalized model for fertility decisions in terms of labor force participation is created. I also discuss potentials for implementation and furthering the model.

Table of Contents

List of Figures

List of Tables

1. Introduction

Many studies have researched the relationship between female labor force participation and fertility rates. Interest in the area began as early as the 1960's. As women began entering the labor force in increasing numbers, many developed nations saw fertility rates drop indicating that there was a possible cause and effect relationship. After 50 years of research, the widespread consensus is that increased female labor force participation does cause decreases in fertility rates. Given this relationship, it is reasonable to assume that the decision-making process behind having a child has changed as women have entered the labor force. With opportunities outside of the home, the cost of having children has increased and women have adjusted their behavior accordingly. In this paper, I will create a model for the fertility decision-making process in terms of a woman's labor force participation.

One of the most important considerations made in the creation of the model will be the dynamic nature of the process. Labor force participation is not a static endeavor. When a woman enters the labor force, her income and experience level are both very low. As she gains experience, her income rises, potential opportunities change, and she becomes a more valuable member of the labor force. A woman participating in the labor force experiences many changes in her income, human capital, and opportunities throughout her career. She gains new information about the opportunity costs of having a child at a very high rate. Due to this constant gathering of information as well as the dynamic nature of the process, a Bayesian model seems most appropriate.

To address the dynamicity of fertility decisions, a Bayesian model will be used. Bayes' Theorem allows the distribution of a parameter to be updated as new information is gathered. This lends itself to the process of fertility decision-making in many ways. Women are constantly gaining information about their opportunity costs and cannot be expected to make a decision as they enter their fertile years that will continue to be applicable as they progress through the labor force. Thus, as they gather information about newly pertinent opportunity costs, women update their decision-making process and come to a new conclusion about whether or not to have children. Therefore, I believe a Bayesian model is most suitable for the process.

I propose that women are implicitly considering their opportunity costs each time they contemplate whether or not to have a child. This model is a formalization of the thought processes each woman in the labor force uses when making fertility decisions. I created this model because I feel that previous models ignore many important facets of the process, particularly the dynamic nature, in favor of simplicity. The goal of this model is to expand upon older models and include the dynamic nature of the issue while not becoming overcomplicated. To do this, Bayes' Theorem in conjunction with economic theories of production and utility will be used to create a dynamic model for the fertility decision-making process in terms of labor force participation.

2. Bayesian Statistics

The creation of Bayesian statistics is nominally attributed to Thomas Bayes, an English minister and mathematician during the 1700's (Press, 1989). However, many point to James Bernoulli as the initiator of the idea of inferential probability. Both men articulated the concept of predicting a probability based on prior knowledge about the occurrence. Bernoulli first expressed the idea in a conceptual framework in his 1713 work, "Ars Conjectandi", while Bayes produced a more mathematical approach in an article in "Philosophical Transactions of the Royal Society" published posthumously in 1763 (Press, 1989). Both articles focused on predicting the probability that something would occur given prior information about the occurrence of similar events. Pierre-Simon Laplace then developed a formal theorem for the concept in 1774 (Press, 1989). Regardless of who initiated the theory behind Bayesian inferential statistics, it has continued to be developed over time and has many uses in today's world.

Most often, Bayesian statistics is used when nothing is known about an unobservable parameter. The lack of information leads to assumptions being made about the distribution of the parameter. Bayesian statistics mediates these assumptions through a process that updates the distribution as new information is collected through sampling. This process not only updates the probability of the event occurring, but also allows for any potential decisions regarding the parameter to be analyzed. The concept of updating the distribution of a parameter is what most obviously differentiates Bayesian statistics from the more commonly used classical statistics.

2.1 Bayes' Theorem

The basic theorem proposed by Bernoulli, Bayes, and Laplace these men that forms the basis for Bayesian statistics is known as Bayes' Theorem. The theorem is used to update the distribution of a parameter given a set of events, as more information about that parameter or the events is gathered. Essentially, the process begins with the researcher making an educated estimation of the distribution of the parameter. Then, they must also determine the likelihood of each event. Using this information, they can update the distribution, making it more and more likely to be representative of the actual distribution the more times Bayes' Theorem is repeated.

The theorem states that the posterior density for a parameter θ is given by

$$h(\theta \mid x_1, x_2, \ldots, x_n) = \frac{L(x_1, x_2, \ldots, x_n \mid \theta) g(\theta)}{\int L(x_1, x_2, \ldots, x_n \mid \theta) g(\theta) d\theta}$$

where the likelihood function is $L(x_1, x_2, \ldots, x_n \mid \theta) = f(x_1 \mid \theta) f(x_2 \mid \theta) \ldots f(x_n \mid \theta)$ such that $X_1, X_2, \ldots, X_n$ are independent, observable random variable vectors with identical distributions and the probability mass/density function is given by $f(x_i \mid \theta)$ for all $i = 1, 2, \ldots, n$, and prior density function $g(\theta)$. This can be further simplified because the denominator depends on the probability mass/density functions of the various x_i's and not on θ. The denominator is actually equal to the marginal distributions of the various X_i's. The theorem can be further simplified to the following,

$$h(\theta \mid x_1, x_2, \ldots, x_n) \propto L(x_1, x_2, \ldots, x_n \mid \theta) g(\theta).$$

This means that the posterior density of θ given the observed random variables, $x_1, x_2, \ldots, x_n$, is proportional to the likelihood of $x_1, x_2, \ldots, x_n$ given θ multiplied by the prior density.

The first step in using Bayes' Theorem is determining the prior distribution of the parameter. When no information is already known, an uninformed prior distribution is used. When using an uninformed prior distribution, it is usually most reasonable to assume that the parameter is uniformly distributed. This is because every possible value for the actual probability is equally as likely. If a researcher has some opinion about what the distribution may be, they determine the distribution given their beliefs. One commonly used distribution is the beta distribution. In this distribution, α and β are shape parameters and, thus, as they are changed, the shape of the distribution changes accordingly. See figure 2.1.1 below to examine how the distribution changes when α and β are changed. However, it is possible to use any distribution for the prior in Bayes' Theorem.

Secondly, the likelihood of the various events given the parameter, θ, must be determined. This is where data comes into play. The likelihood of each event is determined through available data. To determine the likelihood function, each $f(x_i \mid \theta)$'s for all $i = 1,2,\ldots n$ must be ascertained by collecting data. Then, all of the

Fig. 2.1.1

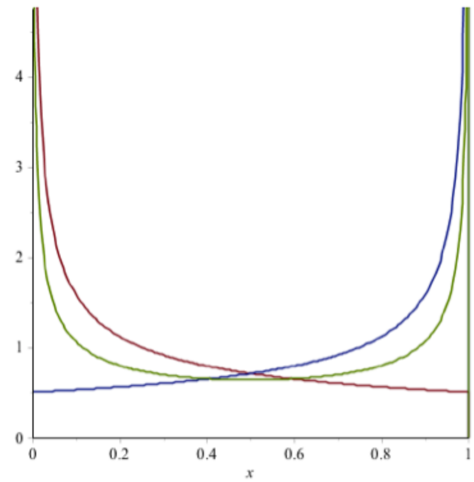

red: α=1 β=0.5
blue: α=0.5 β=1
green: α=0.5 β=0.5

r

$f(x_i\,|\,\theta)$'s are multiplied to create the likelihood function. Once this information has been established, the prior and the likelihood function are multiplied, resulting in a distribution that is proportional to the posterior distribution of θ given various events $x_1, x_2, \ldots, x_n$. To obtain the true posterior distribution, divide by $\int L(x_1, x_2, \ldots, x_n\,|\,\theta)g(\theta)d\theta$, which is a constant with respect to θ but does not depend upon the observed data $x_1, x_2, \ldots, x_n$.

To examine Bayes' Theorem as a whole, consider the following example. Suppose that θ is the parameter of interest and no information is known about the distribution of θ. Thus, it is reasonable to assume an uninformed, uniform prior. Further, let Y be a random variable vector with a binomial distribution. Hence, $f(y\,|\,\theta) = \binom{n}{y}\theta^y(1-\theta)^{n-y}$ where n is the number of independent trials of an experiment and y is viewed as the number of "successes". Therefore, the likelihood function would be $L(y\,|\,\theta) = \binom{n}{y}\theta^y(1-\theta)^{n-y}$. Bayes' Theorem states that the posterior distribution would be given by

$$h(\theta \mid y) = \frac{\binom{n}{y}\theta^{y}(1-\theta)^{n-y}g(\theta)}{\int_0^1 \binom{n}{y}\theta^{y}(1-\theta)^{n-y}g(\theta)d\theta}.$$

However, given that $g(\theta)$ was assumed to be uniform, we know

$$g(\theta) = \begin{cases} 1 \; if \;\; 0 < \theta < 1 \\ 0 \;\; otherwise \end{cases}.$$

So, for example, suppose that in the first test, n = 50 and y = 20. Then the posterior distribution would be as follows,

$$h(\theta \mid y) = \frac{\theta^{20}(1-\theta)^{30} \cdot 1}{\int_0^1 \theta^{20}(1-\theta)^{30} \cdot 1 d\theta}.$$

Thus, the posterior distribution is a beta distribution with $\alpha = 21$ and $\beta = 31$ and a mean of $E(\theta \mid y) = 21/52 \approx 0.404$ and a standard deviation of $\sigma = 0.00454$. This process can be repeated as many times as the researcher wants. Hence, suppose that the researcher decides to repeat the process again and this time does $n = 100$ independent trials with $y = 50$ successes. The posterior distribution would again be updated. However, it is important to note that instead of using $g(\theta)$, the new prior distribution would be the previously determined posterior distribution $h(\theta)$. This results in a new posterior distribution as follows,

$$h_1(\theta \mid y) = \frac{\theta^{50}(1-\theta)^{50}\theta^{20}(1-\theta)^{30}}{\int_0^1 \theta^{50}(1-\theta)^{50}\theta^{20}(1-\theta)^{30} d\theta}$$

$$h_1(\theta \mid y) = \frac{\theta^{70}(1-\theta)^{80}}{\int_0^1 \theta^{70}(1-\theta)^{80} d\theta}.$$

Again, this is a beta distribution with $\alpha = 71$ and $\beta = 81$ and a mean of $E(\theta \mid y) = {}^{71}\!/_{152} \approx 0.467$ and a standard deviation $\sigma = 0.00163$. Note that the distribution remains a beta distribution regardless of how many times the process of updating is repeated. This means that the prior and posterior distribution are conjugate and in this case, the conjugate prior is a beta distribution.

This process of repeating Bayes' Theorem could be continued as many times as deemed necessary, and each time the posterior distribution would become closer to estimating the true distribution of the parameter, as demonstrated by the notable decrease in the standard deviation in the example above. The repetition of the theorem would continue as long as new information could be gathered about the event(s), and, thus, the posterior distribution would become increasingly well informed. Theoretically, a posterior distribution that has been updated using information gathered from the collection of data would be closer to representing the actual distribution of the parameter than the original prior distribution. Once a posterior distribution has been determined, it can be used for many other purposes in Bayesian statistics.

2.2 Bayesian Decision Analysis

Using Bayes' Theorem and a process called Bayesian decision analysis, it is possible to evaluate decision-making. Bayesian decision analysis examines utility in a way similar to economic theory. Each action has a set of associated consequences and each of those consequences is given a value, or utility, based on individual preferences. This utility can be expressed as a function of the consequences in the same way as described above. That is, $U(q_1)$ is the utility function for a particular consequence, q_1. Similarly to economic theory, Bayesian decision analysis assumes that all consequences can be ordered according to preference. Preferential ordering is given a particular notation. For example, if a person does not prefer consequence q_1 to consequence q_2, this is written as $q_1 <\cdot q_2$. Similarly, if a person either does not prefer q_1 to q_2 or is indifferent between the two, it is written as $q_1 \leq\cdot q_2$.

Based on preferences ordering, there are certain properties that a person's utility function $U(q_1)$ must adhere to. The first is that if $q_1 <\cdot q_2$, then $U(q_1) < U(q_2)$. Similarly, if a person is indifferent between the two consequences, their utility functions would be equal, so $U(q_1) = U(q_2)$. Secondly, if $U(q_1) = pU(q_2) + (1-p)U(q_3)$ this implies that a person is indifferent between consequence q_1 and taking a gamble between q_2 and q_3 with probabilities p and $(1-p)$ respectively (Press, 1989).

Often when attempting to measure utility, it is easier to estimate this using a loss function. This loss is the difference between the utility derived from the best possible consequence given the resources and the utility of the consequence that was actually received. To measure this difference, a point estimator is used. This estimator is often assumed to be the mean of the posterior distribution of θ.

Suppose, using an example presented in "Bayesian Statistics, Principles, Models, and Applications", that the loss function is given by $L(\theta,\hat{\theta}) = (\theta - \hat{\theta})^2$ and that the parameter θ is continuous and has a posterior density function given by $h(\theta \mid x_1, x_2, \ldots, x_n)$ (Press, 1989). Then, the expected loss is

$$E[L(\theta,\hat{\theta})] = \int_{-\infty}^{\infty} (\theta - \hat{\theta})^2 h(\theta \mid x_1, x_2, \ldots, x_n) d\theta.$$

To show that this is minimized with respect to $\hat{\theta}$ we take the derivative of the expected loss function and set it equal to zero then solve for $\hat{\theta}$,

$$dE \big/ d\hat{\theta} = \int_{-\infty}^{\infty} -2(\theta - \hat{\theta}) h(\theta \mid x_1, x_2, \ldots, x_n) d\theta = 0.$$

This results in $E(\theta) = \hat{\theta}$. So, the expected value of the parameter θ is equal to the mean of the posterior distribution.

2.3 Hypothesis Testing

In Bayesian statistics, it is possible to test two hypotheses, H_0 and H_1, against one another. When comparing hypotheses, in each hypothesis the parameter θ is equal to one or more specified value. Further, it is assumed that the hypotheses cannot all be true and must cover all possible values of θ. So, if $H_0 : \theta = \theta_0$ and $H_1 : \theta = \theta_1$ and there are no other hypotheses, then θ_0 and θ_1 must be the only two possible values of θ. Each hypothesis also has a probability, P_0 and P_1, respectively associated with it. If no information is known about these probabilities, the default choice is $P_0 = P_1 = 0.5$ (Berger, 2001). Thus, the posterior probability of θ given various random variables under the vectors $X_1, X_2, \ldots, X_n$, is

$$\Pr(H_0 \mid x_1, x_2, \ldots, x_n) = \frac{P_0 \cdot L(x_1, x_2, \ldots, x_n \mid \theta_0)}{P_0 \cdot L(x_1, x_2, \ldots, x_n \mid \theta_0) + P_1 \cdot L(x_1, x_2, \ldots, x_n \mid \theta_1)}.$$

The likelihood function is obtained in the same way as explained previously.

Another possible approach is to use Bayes' Factor, sometimes referred to as the weighted likelihood ratio, or the posterior odds ratio. This method is notationally simpler than the previously outlined method. In a situation in which each hypothesis has equal probability, Bayes' factor is simply equal to the likelihood ratio. In this case, the posterior odds ratio is given by

$$\frac{\Pr(H_0 \mid X_1, X_2, \ldots, X_n)}{\Pr(H_1 \mid X_1, X_2, \ldots, X_n)} = \left[\frac{\Pr(H_0)}{\Pr(H_1)}\right]\left[\frac{\Pr(X_1, X_2, \ldots, X_n \mid H_0)}{\Pr(X_1, X_2, \ldots, X_n \mid H_1)}\right].$$

In a situation in which H_0 is a specified value but H_1 can be multiple values, Bayes' factor is the ratio of the likelihood given H_0 to the average likelihood of H_1. This is represented as

$$\frac{\Pr(H_0 \mid X_1, X_2, \ldots, X_n)}{\Pr(H_1 \mid X_1, X_2, \ldots, X_n)} = \left[\frac{\Pr(H_0)}{\Pr(H_1)}\right]\left[\frac{\Pr(X_1, X_2, \ldots, X_n \mid H_0)}{\int L(X_1, X_2, \ldots, X_n \mid \theta) p(\theta) d\theta}\right]$$

where $p(\theta)$ is the prior density of the data under H_1.

For example, suppose $H_0 : \theta = 1$ and $H_1 : \theta \neq 1$ and there is one random variable vector *Y*. Then, according to the first method, the probability of the first hypothesis H_0, is given by

$$\Pr(H_0 \mid y) = \frac{P_0 \cdot L(y \mid \theta = 1)}{P_0 \cdot L(y \mid \theta = 1) + P_1 \cdot \int L(y \mid \theta \neq 1) p(\theta) d\theta}.$$

Note that there is a slight change in the denominator compared to the method originally described. This is due to the fact that H_1 no longer hypothesizes that θ is equal to a specific value, but rather that it can be many values. Thus, the second half of the denominator represents the average likelihood of data under H_1 instead of simply the likelihood of H_1 in a similar manner as the Bayes' factor method. Similarly, the Bayes' factor method would result in

$$\frac{\Pr(H_0 \mid y)}{\Pr(H_1 \mid y)} = \left[\frac{\Pr(H_0)}{\Pr(H_1)}\right]\left[\frac{\Pr(y \mid H_0)}{\int L(y \mid \theta)p(\theta)d\theta}\right].$$

The Bayesian method of hypothesis testing yields the probability that a hypothesis is true. Thus, in the first method, $\Pr(H_0 \mid x_1, x_2, \ldots, x_n)$ is the probability that the first hypothesis is true given the data provided by $x_1, x_2, \ldots, x_n$. So, if the first method is used for each hypothesis, all of the probabilities can be compared to see which has the greatest probability of being true. Similarly, the second method of using Bayes' factor or the posterior odds ratio can be used as comparison of the probabilities as is already incorporated into the method. Thus, the ratios would be compared to one another if there were more than two hypotheses. If there are only two hypotheses, method two would readily reveal which has the greater probability of being true.

Although the methods and processes discussed above have many similarities to frequentist or classical statistical methods, there are also many differences to consider. Each statistical approach has its merits and drawbacks, and being aware of these ensures that results and limitations are more fully understood.

2.4 Bayesian vs. Classical Statistics

One of the main differences between Bayesian statistics and classical or frequentist statistics is that Bayesian models can be updated according to newly gathered information while frequentist models do not change based on new information. Bayes' theorem for a continuous parameter states that given a likelihood function of $L(x_1, x_2, \ldots, x_n)$ indexed by a continuous parameter θ with a prior density function $g(\theta)$, the posterior density function is

$$h(\theta \mid x_1, x_2, \ldots, x_n) = \frac{L(x_1, x_2, \ldots, x_n \mid \theta) g(\theta)}{\int L(x_1, x_2, \ldots, x_n \mid \theta) g(\theta) d\theta}.$$

The prior density function is determined by any previous experience or information about the parameter θ. That is, $g(\theta)$ is subjective to any beliefs the particular researcher may hold about the distribution of θ. Then, this prior distribution is updated using the likelihood function. The likelihood function represents the likelihood of certain events $x_1, x_2, \ldots, x_n$ given θ. So, essentially, as new information is gathered from the data about the likelihood of certain events, the distribution of θ is updated, resulting in the posterior distribution, $h(\theta \mid x_1, x_2, \ldots, x_n)$.

The existence of a prior distribution developed using beliefs and understanding of what the distribution of a parameter θ may look like is one of the main causes of contention regarding the use of Bayesian statistics. Many proponents of the frequentist approach claim that the use of a prior distribution based on beliefs introduces too much subjectivity into a problem. Nonetheless, subjectivity exists in classical statistics as well. In the frequentist approach, a random sample is assumed to match the distribution of the population as a whole. However, this assumption in itself introduces a level of subjectivity. Furthermore, subjectivity is often built into a problem from the start because a researcher decides to pursue a question due to his or her preconceived notions concerning the results. Despite the seeming subjectivity in the Bayesian method, the assumptions are made evident from the beginning and are often of little importance if the sample size is large enough, as the information gathered diminishes the influence of a prior distribution.

Furthermore, an essential component of Bayesian statistics is the treatment of all unknown parameters as random variables. In the classical approach, all unknown parameters, that are not the main concern of the problem, are treated simply as unknown and unchanging. Frequentist statisticians assume that these parameters are constants. In Bayesian statistics, however, every unknown parameter is treated as a random variable. That is, each unknown parameter is treated as though it is a function instead of a constant. So, each unknown variable is described probabilistically.

The treatment of unknown parameters as random variables plays an important role in Bayes' Theorem. In fact, one of the most important aspects of the theorem is that, as information is updated, the posterior density for the unknown parameter θ is also updated. What is implied by the theorem but is not explicitly stated, is that each event $x_1, x_2, \ldots, x_n$ is also an unknown parameter. So, as information is gathered about each variable, their respective cumulative distribution functions are updated and this results in a new estimation for their likelihood function given θ, which, as discussed previously, results in an updated posterior density. Therefore, the assumption that each unknown parameter is a random variable is an integral part of putting Bayes' Theorem into practice.

Another major difference between Bayesian and classical statistics is in the methods for hypothesis testing. In each of the approaches the basic idea of having two mutually exclusive, exhaustive hypotheses is the same. These two hypotheses are referred to as the null and alternative hypotheses. However, the similarities end there. In frequentist statistics, the next step in hypothesis testing is determining and stating the test statistic, selecting a level of significance, and then computing the observed value, which determines whether or not the null hypothesis should be rejected at the given level of significance using p-values. On the other hand, the Bayesian approach simply calculates the probability that each hypothesis is true and then compares those probabilities to determine which hypothesis is more likely. Given this basic outline of each approach, it is obvious that there are huge differences between the two, particularly in the amount of researcher bias involved in the frequentist approach.

Considering that one of the main complaints against Bayesian statistics is its subjectivity, it is interesting to note how much more subjective hypothesis testing is in classical statistics than in Bayesian statistics. Firstly, the researcher must choose a level of significance, immediately introducing subjectivity into the method. If p-values are used, we must first know the distribution of the data, which is an assumption that is often made without explanation. Although this may seem insignificant, this bias can often be hidden much in the same manner as other facets of frequentist statistics, as previously discussed. In Bayesian statistics, the probabilities of each of the hypotheses are compared against one another, eliminating any researcher bias. Furthermore, researchers using frequentist statistics have the ability to manipulate the outcome of a significance test by simply choosing different tests in order to get the desired result (Babu, 2012). This is simply impossible in Bayesian statistics.

Similar to hypothesis testing is the concept of confidence intervals. In frequentist statistics, confidence intervals are used to determine how good an estimate of an unknown parameter is. To do this it is necessary to decide upon a level of confidence, for example, 95%. Then, the resulting interval would include the true value of the parameter in question 95% of the time if the procedure were to be repeated on multiple random samples of the same size. Again, this method introduces subjectivity in that the researcher must choose a level of confidence. Bayesian statistics has a similar feature called credibility intervals. However, the approach is much simpler. Say, for example, a researcher would like to know what the probability is that a parameter, θ, will fall between a and b given the posterior cumulative distribution function of $(\theta \mid X_1, X_2, \ldots, X_n)$ is $F(\theta)$. Then, the interval statement would be

$$\Pr(a \leq \theta \leq b \mid X_1 X_2, \ldots, X_n) = F(b) - F(a) = c$$

where $0 \le c \le 1$. This simply means that there is a probability of c that the parameter is in the interval $[a,b]$. In this method there is no room for subjectivity beyond that already introduced in Bayes' Theorem.

Overall, both classical and Bayesian statistics have their merits and their limitations. Although it is usually only pointed out in Bayesian statistics, each of the approaches introduces some level of subjectivity. It is often easier to argue the subjectivity in Bayesian statistics due to the fact that it is immediately acknowledged and rationalized, while the subjectivity found in classical statistics is often hidden. The subjectivity in Bayesian statistics is introduced from the very beginning in the creation of a prior based on researcher beliefs and knowledge, while the subjectivity in frequentist statistics is found in the choice of distribution, type of test, and of level of confidence. It is important to recognize the limitations of the method that is being used.

2.5 Use of Bayesian Statistics

When deciding upon the approach to use to model a woman's fertility decisions, Bayesian statistics was appealing primarily because of the ability to update the distribution of the likelihood of having children as new information is gathered. As women enter their fertile years, they have very little concept of what their situation will be three years from now, let alone ten years from now. As they progress through their lives and the labor force, they gain a better understanding of what they can expect their situation to be in the coming years. Furthermore, a woman's preferences for having children can also change as she gains an understanding of all that goes into raising children. Therefore, the basic premise of updating the distribution as information is gathered, seemed most appropriate for the dynamic nature of women's fertility decisions.

After learning more about Bayesian statistics, another aspect of the discipline lent itself to the nature of this project. Despite being one of the main points of contention for critics of Bayesian statistics, the subjectivity built into Bayes' theorem allows for transparency in the creation of the model. The necessity of stating any biases or beliefs about the distribution of the parameter is important in maintaining the integrity of the model. Considering this research is partially focused on the enigmatic topic of preferences as well as many abstract situational variables, making assumptions about these aspects of the model is unavoidable. Thus, making the reasoning behind these beliefs as explicit as possible would benefit the applicability of the model and allow for appropriate changes to be made, if necessary.

3. Economic Approach

In addition to the Bayesian statistical methods described in the previous chapter, it is necessary to have an understanding of the economic concepts at play in the model. This approach will involve examining women's ability to participate in the labor force or produce child services as well as their willingness to do so. The interaction between these two will then determine the combination of labor force participation and child services a woman should produce in order to maximize the utility derived from the activities.

To use this model, some assumptions are made about the woman as a rational economic actor. However, due to the fact that having children is an emotional endeavor, as well as an economic one, some of these assumptions will be relaxed to integrate the emotional factors that are often ignored in basic economic analysis. Firstly, the most notably unrealistic assumption is that the economic actor is all knowing and has perfect knowledge. For this research, this assumption will be ignored. As a woman gathers information about her labor force participation, her likelihood of having a child will change. We are basing out analysis off of the knowledge the woman possesses at the time she is making her decision. Another assumption that will also be ignored is that preferences are unchanging. As a woman gets older, he preferences for having children will, obviously, change. Some assumptions that will be held in this model are that preference ordering is complete, transitive, and independent of irrelevant alternatives.

These assumptions will be used to analyze a woman's production possibility frontier, determined by her labor force opportunity costs, and her utility function. Firstly, we will have to gain an understanding of how her production possibility frontier will change as she progresses through the labor force and her opportunity costs change. Then, her preferences will be examined over time to reveal her fertility decisions at specific times throughout her fertile years.

3.1 Opportunity Cost and the Production Possibility Frontier

Opportunity cost is the cost of an activity in terms of the foregone amount of the next best option. When making fertility decisions, women have to consider the opportunity costs that having a child will engender. Regardless of a woman's labor force participation, these opportunity costs can be quite high. A woman must consider the time and money she will need to dedicate to raising the child. However, when considering labor force participation, other costs must be taken into account. Beyond taking up time and money, having a child may cause loss of promotion opportunities, foregone income, and human capital depletion. Presumably, as a woman begins to consider having children, she would take into account the aspects of her life, such as labor force participation, that would be affected, and weighs the consequences of having a child appropriately. Furthermore, other independent events, such as marriage, may play a role in her decision-making, leading to the woman being more or less likely to choose to have children. This research develops a Bayesian model for predicting the likelihood of having children based on labor force opportunity costs and other situational variables.

It is possible to create a graphical representation of opportunity costs as a curve representing the bundle of two activities that can be achieved using all available resources. This can also be referred to as a production possibility frontier. The production possibility frontier for this research is shown below in figure 3.1.1.

Fig. 3.1.1

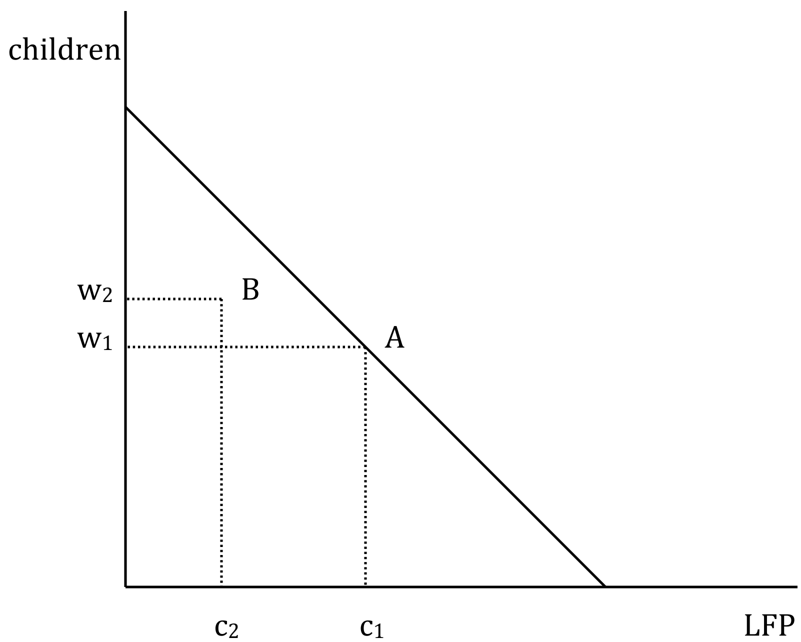

Specifically, a woman can either choose to dedicate her time to having children or dedicate it to her career. Although many researchers measure all opportunity costs by their monetary value and thus group them into one monetary cost, for the purposes of this research it is important to consider each opportunity cost as an individual aspect of a woman's career. In order to obtain bundle B in the example above, a woman would use w_2 amount of time on her career and c_2 time on having children. However, bundle B does not lie on the production possibility frontier and is, therefore, not optimal because it does not use all available resources. A better option, based solely on the production possibility frontier and not on a woman's preferences, would be to spend w_1 on career and c_1 on children to achieve bundle A.

The slope of the production possibility frontier is determined by the marginal rate of transformation. The marginal rate of transformation can be defined as the rate at which the production of one good can be redirected towards producing the other good. The linearity of the production possibility frontier above indicates that the opportunity cost of having a child does not increase with the number of children a woman decides to have. That is, each child a woman decides to have costs the woman the same amount in terms of her labor force opportunity costs. So, a woman experiences the same level of income loss, human capital depletion, and time costs for each child ceteris paribus. The steepness of the curve is determined by the cost of a career relative to the cost of children. A very steep curve would reflect that the cost of participating in the labor force is very high in terms of the amount of child services she would have to give up. Conversely, a flat curve would represent the high career costs of having a child.

It is important to note that the opportunity costs of having a child are dynamic. As a woman gets older and progresses through the labor force, her income, the value of her time, and her potential for human capital development will all change. With these changes, comes a new level of opportunity cost. It is due to these changes that a Bayesian model lends itself to this process; as a woman updates her opportunity costs, she can update her decision-making process accordingly. For example, Gustafsson and Wetzels (1999) proposed that a woman's income, and implicitly the value of her time, throughout her labor force participation begins with a steep upwards curve and begins to level off over time, as seen in figure 3.1.2 below. Therefore, it is important to consider a woman's fertility decision at a given point in time based on the opportunity costs relevant to her at that time. Overall, as a woman's opportunity costs change, so may her decisions regarding her fertility.

Fig. 3.1.2

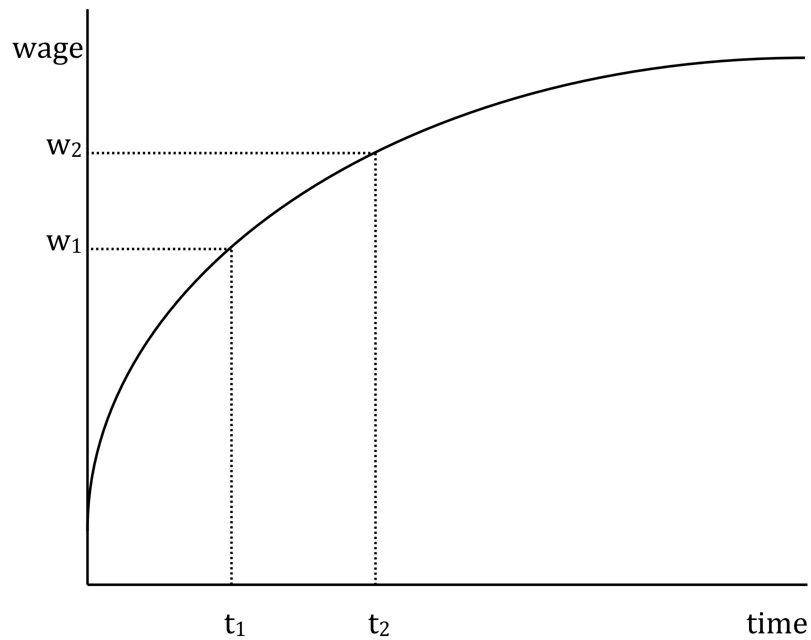

The dynamic nature of the opportunity costs means that the production possibility frontier for having children and participating the labor force will also be dynamic. Thus, figure 1 would be representative a woman's production possibility as a specific point in time rather than over her entire career. So, as a woman gains more information about her labor force opportunities or simply progresses through the labor force, her opportunity costs will change, changing the marginal rate of transformation, and, hence, changing the slope of the production possibility frontier. It is also reasonable to expect the maximum amount of child services or labor force participation to change based on changes in a woman's ability to care for and raise a child monetarily or emotionally as she became older and began to earn more money.

For example, in figure 3.1.3 below, a hypothetical woman's production possibility frontier is given at a specific point in time. The slope of the curve is roughly -1, signifying that it is fairly easy to redirect production in either direction. However, her production possibility frontier after learning that there is the potential for a promotion coming up is very different.

Fig. 3.1.3

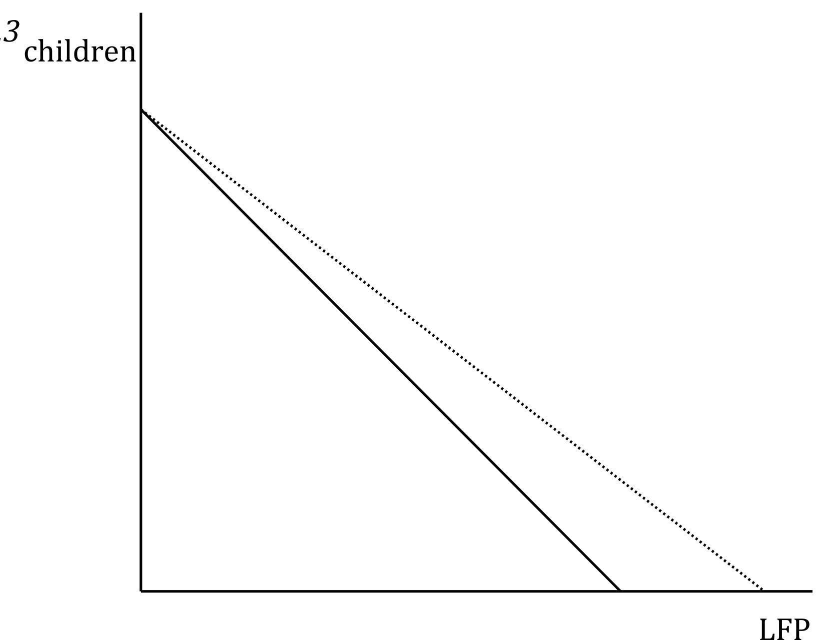

Represented by the dashed line, the slope of her production possibility frontier reflects that the cost of having children in terms of her labor force participation has increased. The slope is now flatter, showing that she would have to give up much more in regards to her career in order to have a child at this time.

Regardless of how informative a production possibility frontier can be, it does not tell the whole story in regards to a woman's fertility decision-making process. To gain a fuller understanding of the amount of resources a woman will direct towards having and raising children and the amount she will direct towards her labor force participation leading to her fertility decision, we must also have knowledge about her preferences.

3.2 Utility and Indifference Curves

According to economic theory, every person has the ability to order a set of bundles in accordance with the level of satisfaction each bundle will yield. This ordering is representative of a person's preferences. Although there is a set of assumptions that economists make about preference ordering, these are not always applicable to actual human behavior. For example, one of the assumptions is that preferences are independent of alternatives. In the theoretical world, when a person has a preference of A over B, the introduction of a third option C will not change the preference of A over B. However, in the real world, a single woman may prefer continuing her career over having children, but when the option of getting married presents itself, her preference for having children may surpass continuing her career. Other assumptions include that more is always better, there is transitivity between options, and all preferences can be ordered.

As previously mentioned, the most common way to analyze preferences is through the use of utility functions. Usually, these function are presented as U(x), which is interpreted at the utility derived from a particular amount of good X. Similarly, these functions can be expanded to include two or more goods in a bundle. This results in functions that take the form U(x_1, x_2,…,x_n) where each x represents a different good. For this research, the utility function would take the form U(c, w) where c is the presence of a child and w is the presence of a career. So, the function U(c, w) would represent the utility a woman derives from having a child and participating in the labor force.

Similar to the production possibility frontier, the slope of the indifference curve reveals important information regarding a woman's fertility decisions. In the case of indifference curves, the slope is determined by the marginal rate of substitution. This marginal rate is the "amount" of labor force participation a woman is willing to give up in order to gain child services. Thus, a steep indifference curve would reveal that a woman is willing to give up a large portion amount of child services, or even forgo having a child, in order to participate more in the labor force. This signifies that she values labor force participation over children. Conversely, a flat indifference curve would represent that a woman values having children more than participating in the labor force.

Utility functions can be combined with the production possibility frontier to determine the optimal bundle of goods to produce. Specifying a value for utility, calculating bundles of goods that yield that utility, and plotting those bundles results in an indifference curve. For different levels of utility, there will be different indifference curves. When plotted on the same set of axes as the production possibility frontier, the optimal bundle of goods can be determined as the tangency point between the production possibility frontier and the indifference curve. This bundle will not only maximize production, as it will inevitably be on the curve, but it will also maximize utility in terms of the production possibility. This is demonstrated in figure 3.2.1 below. According to the woman's preferences, modeled by the indifference curves, and her production possibility frontier, she should have c_1 children and invest w_1 into her career.

Fig. 3.2.1

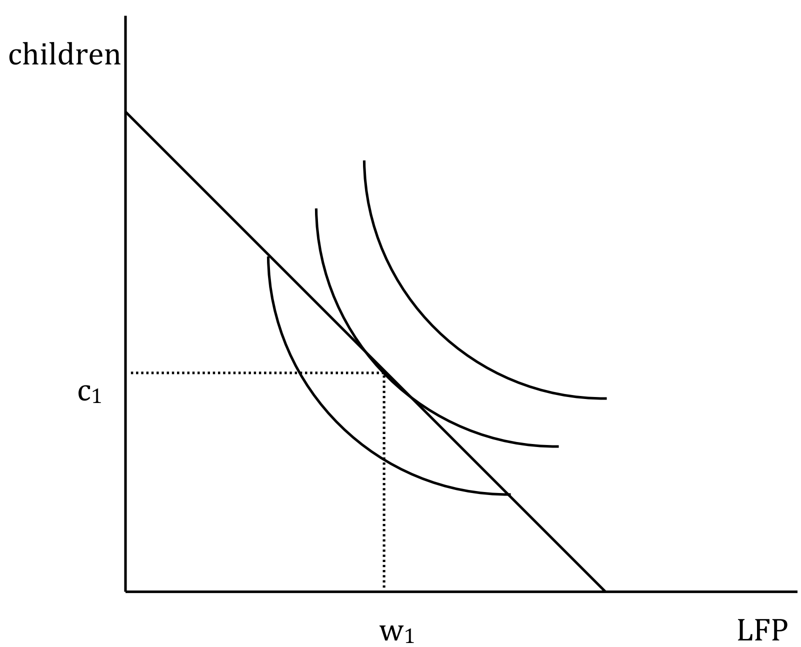

Preferences, and implicitly utility functions, determine how much of her life a woman is willing to dedicate towards labor force participation and having children. By taking the point of tangency between her indifference curve and production possibility frontier, her utility is being maximized in terms of her production capacity. This graphical approach can be translated into a Bayesian model to create a theoretical determination of how women participating in the labor force make fertility decisions.

One of the common assumptions about rational decision-making is that preferences are unchanging. However, this seems to be an unreasonable assumption to make as preferences towards having children often change as a woman gets older. For example, the vast majority of 15-year-old girls do not want to have a child at that point in their lives. However, by the time these women are in their mid-twenties or early-thirties, this preference may have changed. By this point in their lives, societal norms dictate that they should be having children and they are usually more emotionally prepared to raise a child.

This shift is graphically represented below. A solid line represents the hypothetical woman's preferences at 15-years-old and a dashed line represents her preferences at age thirty. Note that the production possibility frontier has not been shifted. This is because a woman at 15 has great potential to participate in the labor force and her potential future income is high compared to her current income, if she has any. A woman at 30 may also still have potential to increase her income and progress in the labor force, an opportunity that may have to be forgone if she were to have children. Thus, the production possibility will be held constant in this example. The point of tangency between the production possibility frontier and her indifference curves can determine the woman's ideal, utility maximizing combination of labor force participation and child services.

Fig. 3.2.2

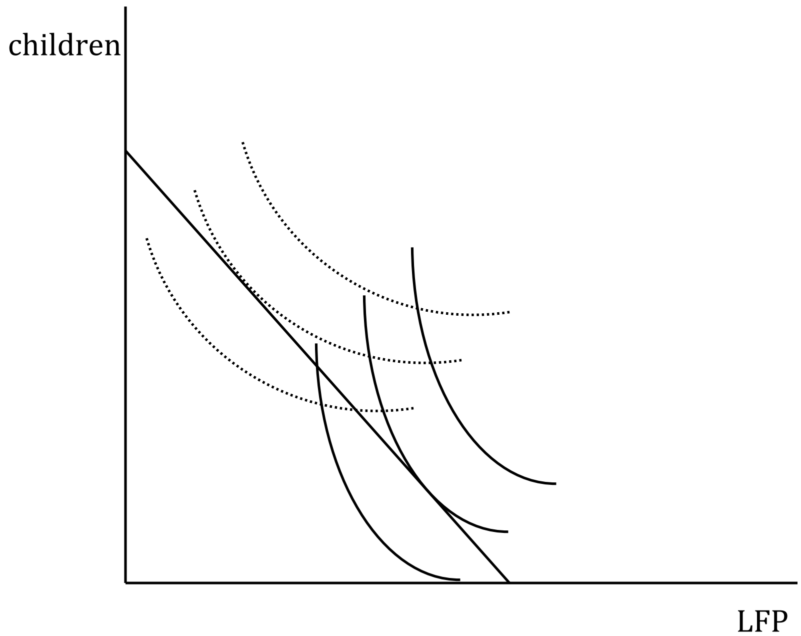

3.3 Conclusions

Based on the assumption that both the production possibility frontier and indifference curves would be dynamic in terms of fertility and labor force participation, Bayesian methods seemed the best fit for this economic model. A woman's opportunity costs as well as several situational variables will affect the likelihood of a woman having a child and this will then affect the utility function for the woman and, thus, her utility maximizing fertility decision.

4. Theoretical Approach

By combining Bayesian statistical methods with the economic concepts of production possibility frontiers and utility functions, a model will be developed for predicting the probability of a woman choosing to have a child given her opportunity costs. Given that a woman's opportunity costs are, as previously mentioned, dynamic, as she gains information about her opportunity costs, the probability of her choosing to have a child can be updated using Bayes' theorem. Her production possibility frontier and utility function can then be updated according to her new likelihood of having a child.

4.1 Interpretation of Bayes' Theorem

When deciding whether or not to have children, a woman takes into account many aspects of her life as well as her preferences at the time she is making this consideration. Thus, the decision changes as the circumstances change. A woman with no husband and low income, but higher potential future earnings and many promotion opportunities, is likely to have a lower probability of having a child than a woman who has been married for five years, has very few promotion opportunities, and whose income has plateaued. This occurs because the opportunity costs of having a child for the first woman are much higher. However, five years down the road, her circumstances may have changed, as well as her preferences, and her probability of having a child may increase. Although a woman may use the same decision process for having children at different ages, her probability of having a child will vary depending upon the information about her opportunity costs and her preferences for having children. In fact, as a woman gains information, she will update her decision-making and the result may change.

For example, a woman at 15 years old, entering what is considered her fertile years, will most likely not want to have a child. Furthermore, the opportunity cost of having a child at 15 is very high. If a woman had a child at this young age, she may be giving up education, which could limit her future income labor force mobility greatly. Most 15-year-olds do not have an income, or at least one that can adequately support a child, nor do they have a reliable career path and thus they know very little about their future income or promotions. So, given this limited information, the woman makes a determination of her probability of having a child. Presumably, at this point, the probability would be very low.

Now suppose this same woman is reconsidering her decision after she has left college. At this point, she is possibly not married, is now entering the labor force and has a relatively low income in comparison to her potential future income, and has many promotion opportunities in the near future. However, she also has a greater ability to care for a child than she did at 15. Furthermore, her preferences for having children may have shifted slightly and having children may be something she considers more appealing than she did at 15. Based on her new information, the woman will update her previous decision. Most likely, the probability of her having a child will increase slightly, but still be low due to the potential for future labor force opportunities and her marital status.

Finally, consider this same woman making a fertility decision at 30. At this point in her career, likely she is married, has moved up in her labor force status, and her future earning potential has begun to level off (Gustaffsson & Wetzels, 1999). She is more adequately prepared to have and raise a child. Additionally, she probably wants children more at this point that she has previously. Thus, she may decide to pursue having children and, as a result, her probability of having a child will be fairly high.

When making fertility decisions, a woman determines whether or not she should have children based upon her prior ideas, her preferences, and any new information that has arisen since she last considered having children. Thus, as a woman progresses through the labor force and, implicitly, her life, she periodically reconsiders having children and the probability of her having a child changes based upon this reconsideration process.

A formal way to model this process of updating her probability of having a child is through Bayes' Theorem,

$$h(\theta \mid x_1, x_2, \ldots, x_n) = \frac{L(x_1, x_2, \ldots, x_n \mid \theta) g(\theta)}{\int L(x_1, x_2, \ldots, x_n \mid \theta) g(\theta) d\theta}.$$

In terms of fertility decisions, θ would represent the probability that a woman would have a child. Thus, Bayes' theorem would give the probability that a woman would have a child given her labor force opportunity costs and situational variables represented by $x_1, x_2, \ldots, x_n$. This probability of having a child will be proportional to the likelihood of the opportunity costs and situational variables given that the woman has a child multiplied by the prior probability of a woman having a child. So, a woman's first decision, the prior distribution, is based on her beliefs about her probability of having a child and her preferences and as she gains information about her opportunity costs and situation, represented by the likelihood function, she can calculate her posterior distribution giving her new probability of having a child. Then, when she wants to reevaluate her decision based on new information, the posterior distribution becomes the new prior and the likelihood function is updated according to the new information. Furthermore, due to the treatment of all unknown parameters as random variables, each opportunity cost variable would be treated probabilistically. That is, each of the opportunity cost variables would be treated as though they could take on various values due to unknown information about each particular cost.

For example, one of the opportunity costs a woman participating in the labor force may have to consider when deciding to have a child is her foregone income. As a woman progresses through the labor force, she may gain information about her future earning potential. This means that she can make a better estimate of her foregone income if she were to have a child than she was previously able to. Therefore, based on this better estimation of that one opportunity cost, the likelihood function of the opportunity costs if she had a child is updated and, when multiplied by the prior density function of having a child, the new posterior density is given based on the information she obtained about her earnings. This is applicable to each of the possible opportunity costs of having a child including time costs, promotion opportunities, and human capital depletion due to time away from work. It also applies to situational variables including whether or not a woman is married and the presence of other children in the family.

4.2 Distribution of Opportunity Costs

One unique feature of Bayesian statistics is the treatment of all variables as random, resulting in every variable having a distribution. In the case of the opportunity cost variables $x_1, x_2, \ldots, x_n$, there are some important considerations to be made when choosing an appropriate distribution. Firstly, the costs must be distributed among negative values. Secondly, it is important to take into account how the opportunity costs should be distributed for different values of the parameter θ.

To address the second consideration, consider how the opportunity costs would change as the parameter becomes larger or smaller. For example, a woman who is unlikely to have a child and, thus, has a parameter $\theta \approx 0$ would have very little chance of incurring any opportunity costs. Therefore, the distribution of the opportunity costs should be focused close to 0. For a woman who is almost certainly going to have a child, parameter $\theta \approx 1$. In this case, the woman would most likely incur a high opportunity cost, represented by a large negative number. This would result in the distribution of opportunity costs being focused near the average value for the opportunity cost of having children. If $\theta \approx 0.5$, a woman has a greater likelihood of having a child and, hence, a greater chance of incurring opportunity costs. This indicates that the distribution of opportunity costs should be greater than 0 for most women but less than the opportunity costs for $\theta \approx 1$. Thus, the numerical value of the opportunity costs when $\theta \approx 1$will be a larger negative number than the costs when $\theta \approx 0.5$. This approach reveals that as the parameter gets closer to one, the mean of the distribution shifts away from 0 as the likelihood of incurring more opportunity cost becomes greater.

Beyond examining the behavior of the mean of the distribution, the variance of the distribution is also an important consideration. Obviously, the numerical value of the opportunity costs of having a child cannot be any higher than 0. Thus, the probability density function cannot exceed 0 for any value of θ. Further, larger values of θ should be associated with a probability density function with greater variance. This occurs because women have some level of choice in the degree to which they experience the opportunity costs of having children. For example, many women will choose to take a brief time off of work, resulting in possible loss of future income, missed promotion opportunities, and human capital depletion. In this case, the woman will experience an average opportunity cost. However, a woman could potentially choose to remove herself from the labor force for 18 years, resulting in a very high opportunity cost, although this would be less likely to happen than a brief maternity leave. Occurring even more rarely, a woman could take almost no time away from work, thus incurring very little opportunity cost. Intuitively, the probability density function should be skewed to the left so that very high levels of opportunity cost are not very likely to occur and extremely low opportunity costs are even less likely to occur.

Given this intuitive picture of the distribution of opportunity costs, the type of distribution that models the description best is the gamma distribution. In order to address the first consideration mentioned above, the distribution must be made negative so that the opportunity costs are distributed across negative values rather than positive ones. To apply the distribution to the opportunity cost of having children there are some important decisions to be made about the shape (α) and scale (β) parameters of the distribution. To create a measurement of opportunity costs that can be easily compared, the scale parameter will be held constant at one. The shape parameter will be a function of the parameter θ. This is appropriate because, as discussed above, the opportunity costs vary according to different likelihoods of having children. The intuitive description of the relationship between θ and the shape of the probability density function, indicates an appropriate function of θ for α to be as simple as $M\theta$. Thus, the distribution can be denoted Gamma($M\theta$, 1).

It is reasonable within the model to suppose that M is representative of the average opportunity cost of having children. Thus, as θ changes, the probability of incurring the full opportunity cost decreases. This assumption also allows for the opportunity cost to be greater than M. This would occur because there is a chance that a woman would incur more than the average opportunity cost. Assume, for simplicities sake, that $M = 50$. Then, using the examples above, let $\theta = 0.5$ and $\theta = 1$.

The probability density functions for the two values of θ are plotted below in figure 4.2.1. Note that, as described, when $\theta = 1$, the probability density function is focused around 50 and when $\theta = 0.5$, the probability density function is focused slightly lower. Overall, this is a fair representation of the intuitive distribution of opportunity costs.

Fig. 4.2.1

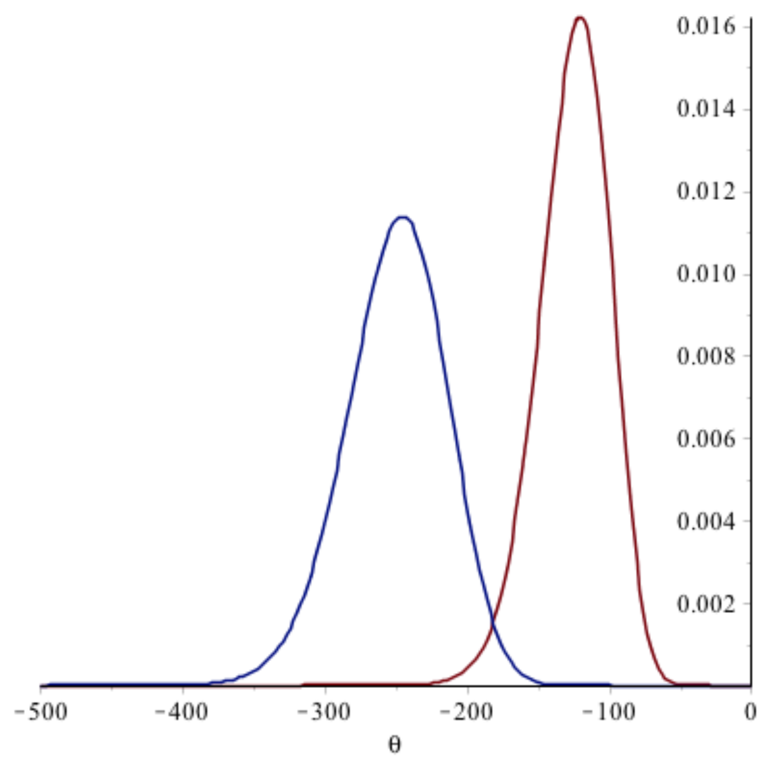

This model for the distribution of labor force opportunity costs for women deciding to have children can be used with Bayes' Theorem to predict the distribution for the likelihood of having children.

4.3 The Likelihood of Having Children and the Utility Function

As previously discussed, one of the main components in a woman's decision-making process in regards to her fertility is the utility she will derive from having a child. Thus, the likelihood of having a child, θ, interacts with a woman's utility function. More directly, θ affects the amount of work and child services a woman is likely to produce. Therefore, it is reasonable to assume that both child services and work are functions of θ. Recall that the utility function for deciding to have children is denoted as U(c, w) where c and w represents a bundle of work and child services. Hence, assuming that both child service and work are functions of θ, the utility function can be rewritten as $U(c(\theta),w(\theta))$. This assumption allows us to make predictions about how the utility function will change as the distribution of θ is updated using Bayes' Theorem.

First, assume that the posterior distribution of θ reveals that the probability that the likelihood of having a child is close to one is very high. This implies that woman is likely to have a child soon. It is reasonable to assume, then, that the utility function for that woman would reflect that the utility derived from having a child is greater than the utility she would derive from work. As discussed in the section on the economic approach, deriving more utility from child services results in a flat indifference curve because a woman is willing to give up a large amount of work in order to gain the ability to produce more child services. In other words, the marginal rate of substitution is low. A depiction of this can be seen in figure 4.3.1 below.

Fig. 4.3.1

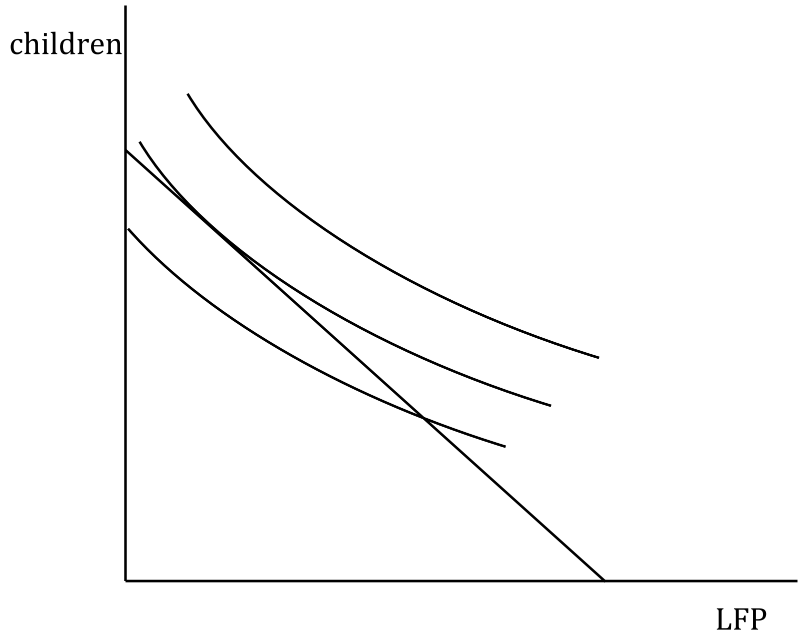

Note that the point of tangency with the production possibility frontier is at a combination at which the amount of child services is high and the amount of work is low.

On the other hand, assume that the posterior distribution shows a very high probability of a woman being unlikely to have a child. This would mean that the woman is very unlikely to have a child. Conversely to the previous example, we can assume the woman would derive more utility from working than having a child. This would result in the marginal rate of substitution being high and the slope of the indifference curve, therefore, being steep. Thus, a woman would be willing to give up large amounts of child services in order to gain the ability to work. Figure 4.3.2 below shows a depiction of this example.

Fig. 4.3.2

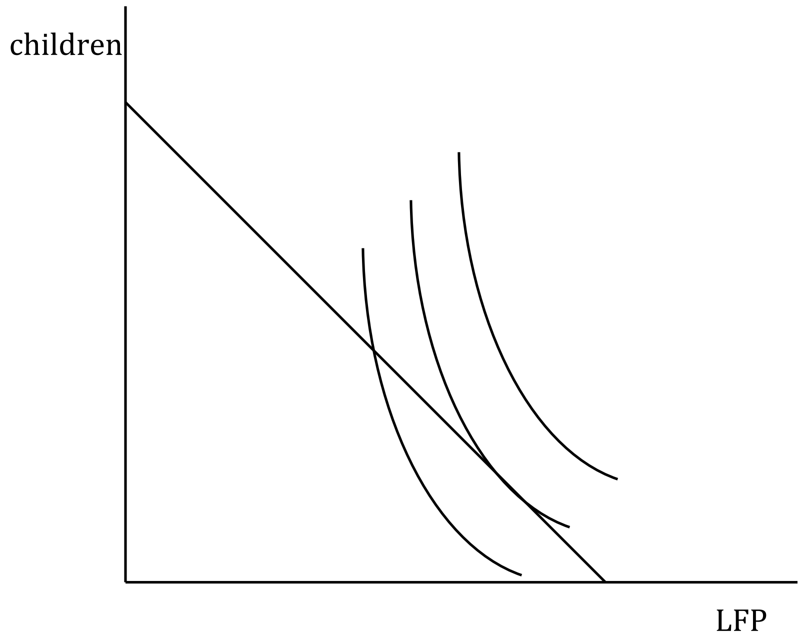

In this depiction, the point of tangency is very different from in figure 1. The amount of work a woman will produce to maximize her utility is very high and she will dedicate almost no resources to producing child services.

Given these examples, it is simple to deduce the relationship between the parameter θ and the utility derived from child services and work. When θ is high, a woman is likely to want to have a child and thus would have a flat indifference curve. When θ is low, a woman most likely would prefer to continue working and, thus, have a steep indifference curve. It is reasonable to expect that woman's likelihood of having a child affects how she views the utility derived from having a child. However, this does not depict the whole story. A woman's production possibility frontier is also going to change according to her likelihood of having a child.

4.4 The Likelihood of Having a Child and the Production Possibility Frontier

Similarly to the utility function, a woman's production possibility frontier will also be affected by her likelihood of having a child. We continue to assume that both child services and work are functions of θ. Therefore, as a woman updates her distribution of θ, we expect to see a resulting shift in the production possibility frontier. This creates a cyclical change so that the distribution of θ affects the production possibility frontier, which signifies a change in the opportunity costs associated with having a child, which, in turn, affects the distribution of θ given the opportunity costs and situational variables.

For example, assume that the posterior distribution of θ signifies that the probability that the likelihood of having a child is close to one is very high. Thus, a woman is very likely to have a child. Not only will the change her utility function, but it will also affect her production possibility frontier. If a woman has a child, she must produce some level of child services. Hence, we can expect there to be a minimum amount of child services produced. This also limits her ability to participate in the labor force leading to a downward shift in the maximum amount of labor force participation. Additionally, we can expect the slope of the production possibility frontier to change due to changes in the opportunity costs of having more children. The slope may become steeper or flatter depending on the information gathered by the particular woman. In this example, we assume that the opportunity cost of having another child in terms of labor force participation are lower, resulting in a steeper production possibility frontier. This can all be seen in figure 4.4.1 below.

Fig. 4.4.1

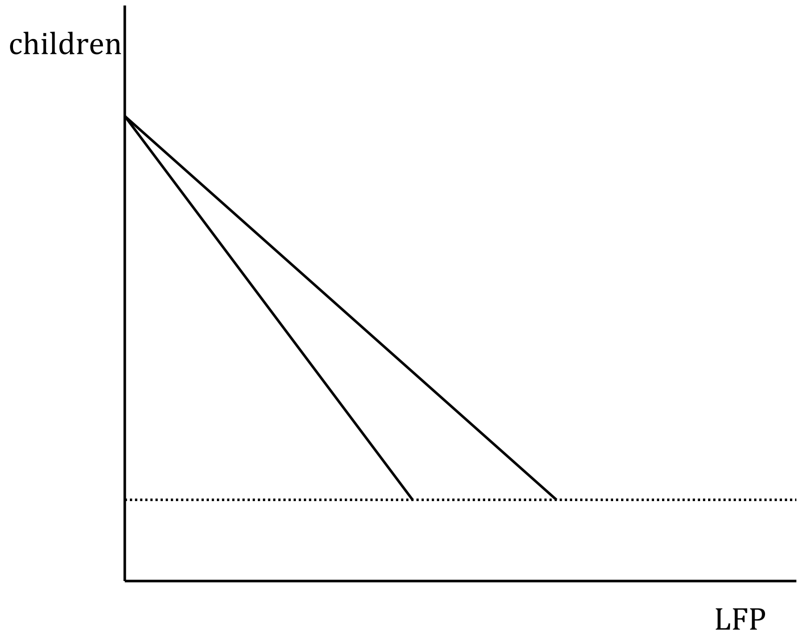

Conversely, consider a distribution of θ that reveals the probability that the likelihood of having a child is zero is very high. This signifies that a woman is not likely to have a child. Considering that no child will be present, there is no minimum amount of child services that will have to be produced. Further, continuing to participate in the labor force means that a woman may receive a raise or promotion resulting in the ability to produce more labor force participation and children. Thus, the maximum amount of each that can be produced will increase. Assume, in this case, that the opportunity costs stay the same so that the slope of the production possibility frontier does not change. This is shown in figure 4.4.2 below.

Fig. 4.4.2

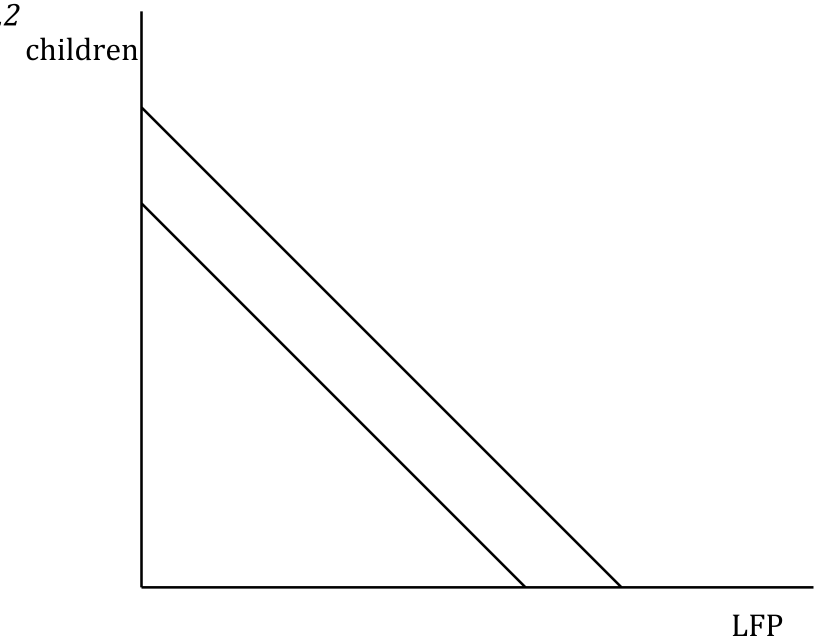

These changes alone do not tell the whole story of a woman's fertility decisions. It is necessary to put all of the pieces together in order to create a full picture of how Bayesian statistics can model the decision-making process.

4.5 The Model Completed

Given the distributions of the opportunity costs discussed above, it is possible to use Bayes' Theorem to complete the Bayesian model of fertility decision-making. Recall that we are assuming that the prior distribution for the likelihood having children is uniform. By assuming a uniform prior distribution between zero and one, we are indicating little is known about the likelihood of having children for an individual woman at a given age. In other words, this is an uninformed prior distribution. Further, for reasons discussed in the previous section, the opportunity costs, represented by random variables $X_1, X_2, \ldots, X_n$, are assumed to have a modified gamma distribution. Now, Bayes' Theorem can be used to produce the posterior distribution for the likelihood of having children given the various labor force opportunity costs.

Using the assumed information, we can write Bayes' Theorem for a particular woman's fertility decision as follows,

$$h(\theta \mid x_1, x_2, \ldots, x_n) = \frac{\left(\frac{(-x_1)^{M\theta-1}}{M\theta \cdot \Gamma(M\theta)}\right)\left(\frac{(-x_2)^{M\theta-1}}{M\theta \cdot \Gamma(M\theta)}\right)\cdots\left(\frac{(-x_n)^{M\theta-1}}{M\theta \cdot \Gamma(M\theta)}\right)}{\int_0^1 \left(\frac{(-x_1)^{M\theta-1}}{M\theta \cdot \Gamma(M\theta)}\right)\left(\frac{(-x_2)^{M\theta-1}}{M\theta \cdot \Gamma(M\theta)}\right)\cdots\left(\frac{(-x_n)^{M\theta-1}}{M\theta \cdot \Gamma(M\theta)}\right) d\theta}.$$

Note that the prior distribution is equal to one and, therefore, does not effect Bayes' Theorem in the first instance. This can be further simplified to the following,

$$h(\theta \mid x_1, x_2, \ldots, x_n) = c(x_1, x_2, \ldots, x_n) \cdot \frac{((-1)^n (x_1 \cdot x_2 \cdot \ldots \cdot x_n))^{M\theta-1}}{(M\theta \cdot \Gamma(M\theta))^n}.$$

The constant, c, is the reciprocal of the value of the integral in the denominator.

Now, assume that the woman in question gains more information about her opportunity costs, represented by random variables $Y_1, Y_2, \ldots, Y_k$. Therefore, the new prior distribution would be $h(\theta \mid x_1, x_2, \ldots, x_n)$ given above. Thus, Bayes' Theorem can be used as follows,

$$h(\theta \mid x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_k) = C \cdot \frac{((-1)^{n+k} (x_1 \cdot x_2 \cdot \ldots \cdot x_n \cdot y_1 \cdot y_2 \cdot \ldots \cdot y_k))^{M\theta-1}}{(M\theta \cdot \Gamma(M\theta))^{n+k}}.$$

In this case, C is equal to the constant determined by the integral in the denominator of Bayes' Theorem multiplied by the opportunity costs represented by the variables $x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_k$.

To simplify this posterior distribution equation further, let $r = (-1)^{n+k}(x_1 \cdot x_2 \cdot \ldots \cdot x_n \cdot y_1 \cdot y_2 \cdot \ldots \cdot y_k)$ when the opportunity costs are less than the average opportunity cost M and $r = 1/(-1)^{n+k}(x_1 \cdot x_2 \cdot \ldots \cdot x_n \cdot y_1 \cdot y_2 \cdot \ldots \cdot y_k)$ when the opportunity costs are more than the average opportunity cost M. This yields the conjugate prior,

$$A \frac{r^{M\theta - 1}}{(M\theta \cdot \Gamma(M\theta))^k},$$

such that $r > 0$, $M > 0$, and k is a positive integer representative of the number of opportunity costs collected overall and A is a constant. Thus, gathering the information about the opportunity costs and plugging the data collected into the following equation yields the posterior distribution,

$$h(\theta \mid x_1, x_2, \ldots, x_k) = \frac{\dfrac{(r)^{M\theta - 1}}{(M\theta \cdot \Gamma(M\theta))^k}}{\displaystyle\int_0^1 \frac{(r)^{M\theta - 1}}{(M\theta \cdot \Gamma(M\theta))^k} \, d\theta}.$$

Given this equation for the posterior distribution, it is possible to use examples to explore what various values for each of the variable would yield. To examine the effects of changing to number of opportunity costs, let $r = 1$ and $M = 1$.

Figure 4.5.1 below shows the posterior distribution of θ when there is only one opportunity cost known.

Fig. 4.5.*1*

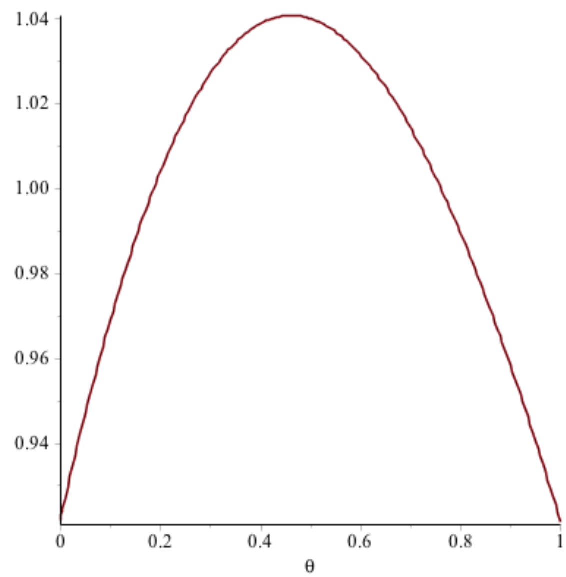

Because very little information is known, the distribution of the likelihood of having a child is fairly spread out. Suppose that 100 opportunity costs are known instead. It would be reasonable to assume that the distribution of the likelihood of having a child is more concentrated considering because the woman making the decision would have more to base her decision on. The plot for this is shown in figure 4.5.2 below.

Now, to investigate the effects of the variable r, assume instead that $k = 1$ and $M = 1$. The distribution for when $r = 1$ can be seen above in figure 4.5.1. Suppose instead that $r = 1/100$. Figure 4.5.3 below shows the plot of the posterior distribution.

Fig. 4.5.2

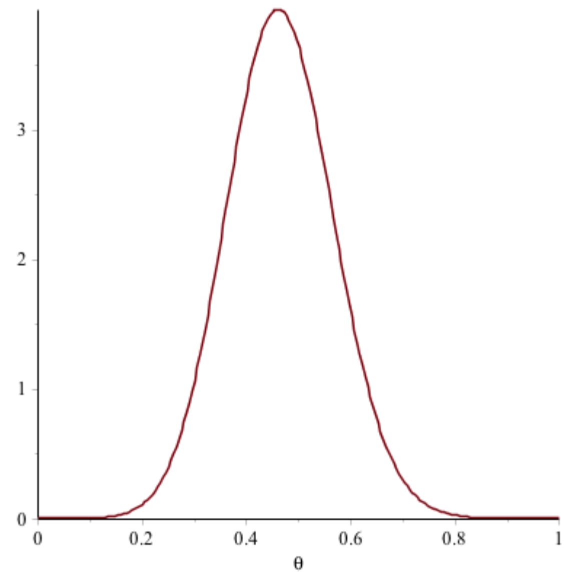

Fig. 4.5.3

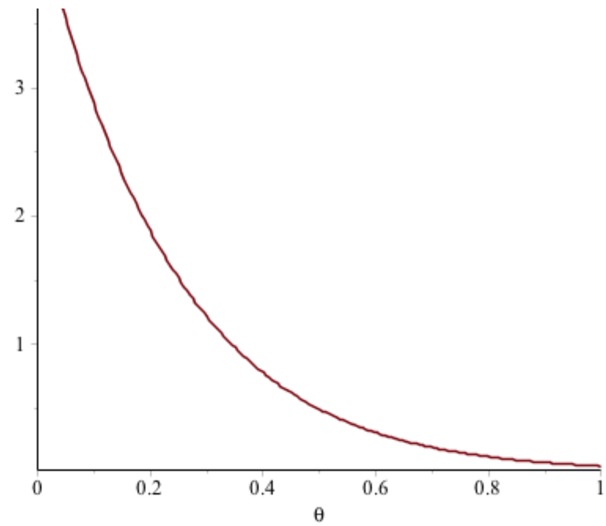

As we can see in the graph above, when the opportunity costs are high, r is low and the likelihood of having a child is distributed among the lower values. Similarly, when the opportunity costs are low, the likelihood of having a child is distributed among high values. As seen in figure 4.5.4 below when $r = 100$. This fits the intuitive picture of how a woman should respond to various levels of opportunity cost.

Now it is also possible that multiple opportunity costs are known and they are high. Suppose, for example, that the opportunity costs are high so that $r = 1/100$

Fig. 4.5.4

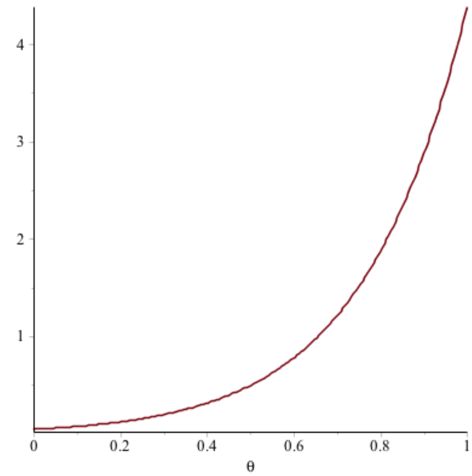

and ten opportunity costs are known. Figure 4.5.5 shows the changes in the probability distribution function.

Fig. 4.5.5

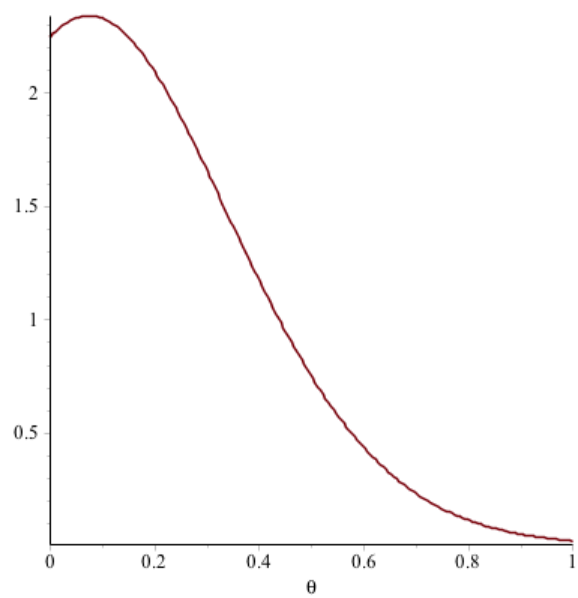

4.6 Conclusions

When examining the entire process, we get a clear picture of how a woman makes her fertility decisions in terms of her labor force participation opportunity costs. As a woman gains more information about her opportunity costs, she has better knowledge and the likelihood of her having a child becomes more defined, as demonstrated by the distribution in figure 4.5.2 when compared to figure 4.5.1. further, as she gains knowledge about the value of those opportunity costs, her likelihood of having a child is shifted to lower probabilities when the costs are high and higher probabilities when the costs are low. This process is repeated whenever the woman gains more information about her labor force opportunity costs to determine a new distribution for the likelihood of having a child.

5. Literature Review

There has been extensive investigation and literature serving as groundwork for this particular research question and methodology. Beginning in the mid-twentieth century, a group of economists turned their attention towards analyzing social decisions rather than market decisions. Many of these economists, most notably, Gary Becker, focused much of their time on investigating decisions made within the home. One of the main topics of research was the rational decision-making process behind producing children. At this time, social norms still dictated that women stay home and take care of the children while men work, but as this changed and women's labor force participation increased, the focus of the research shifted. It was no longer assumed that the woman would stay home with the children, and, thus, the role of female labor force participation in fertility decisions became an important research topic.

The relationship between women's careers and their fertility decisions is still studied to this day, although the methods and research questions have become increasingly more in-depth. For example, in early work the consideration of opportunity cost focused mainly on the cost of time in terms of the wage that could be earned. However, in new research, the greater effect that having children can have on labor force participation and the ensuing opportunity costs has been integrated into the decision-making process. Furthermore, newer research has taken the economic approach beyond the simple supply and demand analysis of the 1960s and expanded it to include more complex methods.

In terms of the application of Bayesian statistical methods to fertility issues on an individual level, research has been limited. However, this could be due to the fact that obtaining enough data to make the research viable is still difficult. Nonetheless, the research that has been done is promising. Researchers have begun applying Bayesian methods, usually reserved for scientific inquiry, to social issues including fertility and labor force participation. The use of Bayesian statistics in social applications illuminates the possibility for further research concerning the relationship between fertility decisions and female labor force participation.

The research discussed in this section covers many facets of the investigation into the relationship between female labor force participation and fertility decisions. Beginning with Gary Becker's work concerning the production of children in the 1970's to a more recent work using Bayesian methods to examine the perceived paradox between women's employment and fertility in European countries, the literature reviewed here spans time and topics. Serving as an analysis of how the research about this subject has changed, this literature review evidences the applicability of this research question.

5.1 "Interaction Between Quantity and Quality of Children" (Becker & Lewis, 1974)

In the early 1960s researchers began examining fertility decisions using an economic approach. The economist that brought the economics of fertility to the forefront was Gary S. Becker with his 1960 article, “An Economic Analysis of Fertility”. His work influenced many other researchers to continue to develop his theories on the interaction between the number of children a woman decides to have and the quality of those children. It also inspired researchers to examine other aspects of fertility, including its relationship with female labor force participation. As an extension of his original paper on fertility, Becker and H. Gregg Lewis wrote “Interaction between Quantity and Quality of Children” (1974).

This article expands upon Becker’s basic explanation of the association between the quantity and quality of children to illustrate that the two are more dependent on each other than any random two household commodities. The main assumption in their proposed model is that the opportunity cost of having children is greater for higher quality children. That is, if a woman decides to have more children, raising a higher quality child will be more expensive in terms of opportunity cost than raising a lower quality child. Becker previously explained that the quality of a child is not dependent on the child’s moral character, but rather the amount of money the parents have designated towards raising the child at a higher caliber (Becker, 1960). To demonstrate the close relationship, Becker and Lewis develop a utility function and budget constraint to determine the equilibrium conditions and examine the implications for income effects and price effects.

The utility function developed in the article depends upon three variables; the number of children (n), the quality of children (q), and the rate of consumption of other household commodities (y). This yields the utility function $U = (n, q, y)$. The budget constraint is dependent upon the price of raising children (p) and the price of household commodities (p_y). This gives the very simple budget constraint $I = nqp + yp_y$. These functions result in a set of equilibrium conditions that must be satisfied,

$$MU_n = \lambda qp = \lambda m_n;\ MU_n = \lambda np = \lambda m_q;\ MU_n = \lambda p_y = \lambda m_y,$$

where the m's are the shadow prices. The important point to take away from the equilibrium conditions is that the opportunity cost of the number of children is positively related to the quality and, similarly, the opportunity cost of the quality of children is positively related to the quantity. So, for example, an increase in the quality of children is more expensive if there are more children because the increase applies to more "units".

When evaluating income effects, Becker and Lewis approach the concept in the traditional manner by changing income while holding all prices constant. They assume that the true income elasticity of the quality of children is larger than the income elasticity of the quantity of children. This is due to the fact that the observed elasticity for the number of children can be negative for very low incomes and then become positive as income increases. This produces an interesting relationship where lower income households are having more children than their wealthier peers. However, Becker and Lewis point out that the true elasticity is never negative, but would still be lower than that of child quality. So, at lower incomes, potential parents demand a certain level of quantity and quality and they make up what they cannot achieve in quality by increasing the number of children they intend to have.

Finally, Becker and Lewis examine the price effects for both quality and quantity of children. They begin this analysis by generalizing the budget constraint to result in the equation $I = np_n + nqp + qp_q + yp_y$, giving the following opportunity costs: $o_n = p_n + qp$, $o_q = p_q + np$, and $o_y = p_y$. Suppose that p_q increases. This would result in an increase of the opportunity cost of the quality of children relative to the other two opportunity costs and would cause the quality of children to decrease. However, this change in the quality of children would have an effect on the opportunity cost of the number of children, and the number of children would increase. Overall, an increase in the price of the quality of children causes a decrease in quality and an increase in quantity. Furthermore, they assume that if all prices are decreased by the same percentage, the quantity of children will tend to increase substantially more than the quality of children.

This article is significant not only because it was written by an economist that sparked the widespread interest in examining fertility through an economic approach, but also because it emphasizes the point that parents, and especially women, are making a decision not only about the number of children to have, but also the quality at which they intend to raise their children. Many researchers have come to refer to this concept of the combination of quality and quantity as child services. This view of quantity and quality being interrelated and somewhat interchangeable, allows for many assumptions about the demand for children or child services. For example, as income increases, there is rarely a large increase in the number of children that a household demands, in fact, some data shows that lower incomes are actually associated with a greater number of children (Becker, 1960). This relationship points to children being an inferior good, which is counterintuitive. However, if child quality is considered, as it is in overall child services, generally, wealthier households spend more money on their children resulting in higher quality children and therefore the demand for child services actually rises with income, making child services a normal good.

Becker's early work on analyzing fertility using economics led to a number of other researchers investigating related topics. One of the main areas of interest that Becker's work inspired was the relationship between female labor force participation and fertility decisions. Many economists noticed the correlation between the increase in female labor force participation in developed countries and the decrease in fertility. This realization lead to a number of articles investigating the extent to which the two were related and then determining the nature of the causal relationship.

5.2 "Intended Childbearing and Labor Force Participation of Young Women: Insights from Nonrecursive Models" (Waite & Stolzenberg, 1976).

The attempts to establish relationships between fertility and female labor force participation began almost immediately after the introduction of fertility as an economic concept. One such article is "Intended Childbearing and Labor Force Participation of Young Women: Insights from Nonrecursive Models" (Waite & Stolzenberg, 1976). Through a sociological approach, the authors intend to establish a causal relationship between fertility and labor force participation. To define the causality, the authors suggest three possible relationships between childbearing and career choices. The first is that fertility intentions determine labor force plans, secondly, labor force plans may determine whether or not a woman has children, or, finally, there may be no relationship between a woman's fertility plans and her career intentions. The goal of this article is to develop a model that allows for the possibility of the three different relationships and to compare the explanatory power of each. This individualistic approach towards a woman's intentions regarding her fertility and labor force participation fits into the framework of fertility decision-making.

The model that Waite and Stolzenberg create focuses on simultaneous equations for labor force participation and fertility expectations. In fact, the two equations are very similar. Both equations include age, marital status, education, and race as explanatory variables. These four variables are included in the model because studies show that they have an effect on both labor force plans and fertility intentions. The equation for labor force participation includes work attitudes and whether or not the woman's mother participated in the labor force as explanatory variables, while the equation for fertility expectations includes ideal family size and the number of siblings of the woman. Finally, labor force participation and fertility expectations are allowed to have direct effects on each other in the model.

To test this model, the authors use data from the National Longitudinal Study of the Labor Market Expectations of Young Women. From 1968 to 1973, The Ohio State University and the U.S. Bureau of the Census annually interviewed 5000 women between 14 and 24-years-old about their plans for the future in regards to career and fertility. By 1973, 91 percent of the interviewees had responded all six times. Participants were asked questions about their labor force plans at the age of 35, how many children they intended to have, and other questions related to the variables included in the authors' model. The authors used responses to the question to run a two stage least squares regression and an ordinary least squares regression to calculate the coefficients for variables including, fertility expectations, labor force participation plans, age, freewill, and the perceived costs and benefits of working.

The authors conclude that fertility intentions and labor force participation are interdependent. Therefore, their relationship is not spurious. The results of the regression analysis, seen in table 5.2.1 below, show that a woman who intends to participate in the labor force tends to plan to have fewer children than her peers who do not intend to work. In fact, according the results of this study, women who intend to work have 0.767 fewer children than women who would prefer to stay at home. This reinforces the notion that female labor force participation has a negative effect on fertility decisions. When estimating their model, they found that the most likely causal relationship was that plans to participate in the labor force affected the number of children a woman intended to have.

Table 5.2.1

	Labor Force Plans as Dependent Variable		Fertility Expectations as Dependent Variable	
	TSLS	OLS	TSLS	OLS
work attitude	-0.16	-0.17	-	-
age	0.07	0.07	0.06	0.04
labor force plans	-	-	-0.32	-0.02
fertility expectations	-0.08	-0.04	-	-
R^2	0.07	0.07	0.26	0.34

These results support the idea that labor force participation does have an effect on fertility decisions. Further, they emphasize the fact that a woman's intentions towards participating in the labor force can lead to her deciding to have less children. Therefore, it is reasonable to conclude that having children is seen as some sort of barrier to labor force participation. My research operationalizes this idea in its use of opportunity costs. If women are letting their labor force intentions affect the number of children they want to have, they must be examining the effect that having children will have on their careers.

Due to the fact that this was published in a sociological journal and not an economic journal, there is no explicit discussion of opportunity costs although the variable "work attitude" measures the perceived costs and benefits of working. However, it is important to note that the authors determined this variable should have no effect on fertility decisions because a relative can look after children while the mother works. This argument would be invalid now because many women do not live near relatives and, therefore, cannot use them at their disposal. Despite this dismissal of the variable, their reasoning implies that women are taking into account the time cost of having children in terms of their labor force participation. In spite of a lack of economic discussion of opportunity costs, this article implies that women as young as 14 are taking into consideration what they must give up in order to have children and are using this to inform the plans for their future.

This article reinforces two foundational aspects of research concerning a woman's fertility decisions in terms of her labor force participation. The first is that women as young as 14 are considering the opportunity costs of having children. Although this is not explicitly discussed within the article, as discussed above, certain variables imply that this is the case. Secondly, and most importantly, this article shows that intentions to participate in the labor force have an effect on fertility plans. The basis of this research is that a woman makes her fertility decisions according to the labor force opportunity cost of having children..

This article was an early attempt at examining the relationship between a woman's intentions, and implicitly her preferences, for having children and her labor force participation. One of the main concerns with the research presented in the article is that it is a point-in-time analysis. This type of static model does not account for the inevitability of changing preferences as a woman gets older. My dynamic model improves upon this point. Furthermore, although it does effectively establish a causal relationship, this article considers the relationship between fertility intentions and career intentions and fertility behavior and career behavior separately and in terms of whether intentions were met. It does little to evaluate the actual interaction between fertility intentions, career importance, hours worked, and actual fertility behavior. However, more recent articles have expanded upon this research to include a more comprehensive view of the relationship.

5.3 "Fertility Intentions, Career Considerations, and Subsequent Births: The Moderating Effects of Women's Work Hours" (Shreffler & Johnson, 2013).

Karina M. Shreffler and David R. Johnson are two economists who have studied the complexities of the relationship between intentions and behaviors in terms of labor force participation and fertility at a more in-depth level. Their article, "Fertility Intentions, Career Considerations, and Subsequent Births: The Moderating Effects of Women's Work Hours", explores the interaction between fertility intentions, career importance, number of hours worked, and fertility behavior. Their main goal was to evaluate the importance of preferences in a woman's fertility behavior and number of work hours. To categorize a woman's preferences, Shreffler and Johnson established three groups that women could belong to. These included work-centered women who placed more importance on career than motherhood, home-centered women who believe motherhood was more important than labor force participation, and, finally, adaptive women who wanted a combination of family and work in their lives. They examine all the relationships to determine what effect each has on whether or not a woman has a child and develop nine hypotheses concerning the various connections.

To test their various hypotheses, Shreffler and Johnson used data from the National Survey of Families and Households. The survey was conducted in two waves. The first wave was in 1987-1988 and included approximately 13000 participants 19 and older. Participants were asked a series of questions about their childhood, marital history, fertility, education, and employment. In 1992-1994, the second wave was conducted with 10000 of the original participants. In this wave the respondents were asked about details of their lives since the first wave. In order to gain an appropriate sample for their research, Shreffler and Johnson eliminated all participants over the age of 39 and those who were non-surgically sterile. After the various restrictions, the sample included 2411 women.

They then developed three models to test their hypotheses. The dependent variable in all models was the existence of a birth between waves. The number of births was not important to the study, and, therefore, only the first birth was significant. The independent variables were work hours, career importance, and fertility intentions, which were scored between one and seven. There were also various control variables including age, education, race, income, marital status, and number of children. They split the analysis into two sections. The first examined the effects of preference on work hours. The second was intended to evaluate the effects of work hours, career importance, and fertility intentions on whether a woman had a child. This was then evaluated using three models. In the first model they only included the control variables and work hours. The second model included all variables from model one and added fertility intentions and career importance. Finally, the third model examined significant interaction and curvilinear effects. They used OLS multiple regression to evaluate the relationships.

Examining the effects of preferences on work hours revealed information concerning their first two hypotheses. They believed that a woman who places a higher value on career would work more hours and that women who had greater fertility intentions would work fewer hours. The first hypothesis was supported. Women who thought their careers were very important worked more hours than their peers. However, the second hypothesis was not supported. There was no significant relationship between the number of hours worked and a woman's fertility intentions.

Their other hypotheses were evaluated through the three-model examination of the effects of work hours, fertility intentions, and career importance on whether a woman would have a child. The first model showed that increased work hours did result in a decrease in the probability of giving birth. So, women who worked more were less likely to give birth. The second model included both fertility and career preferences as well as work hours. This model showed that placing higher value on career resulted in a woman being less likely to give birth and placing higher value on fertility resulted in a woman being more likely to give birth. The third, and final, model was used to evaluate the moderating effects of work hours on the relationship between preferences and giving birth. This model found that home-centered women were more likely to meet their fertility intentions regardless of the number of hours worked. It also showed that women who placed high value on their careers and were working more hours were less likely to have a child.

Shreffler and Johnson's article differentiates itself from other literature because it establishes that there is no association between fertility intentions and the number of hours that a woman works. Interestingly, the research finds that women who worked more hours were actually more likely to meet their fertility intentions than their peers who worked fewer hours. This is significant because it demonstrates the importance of considering preferences in research about fertility. Their results indicate that, while other research places high importance on the number of hours worked and how that measure of labor force participation is leading to the decrease in fertility, it may be more important to examine a woman's intentions.

This article reinforces the premise of this research. It emphasizes the concept that work hours are not the only determining factor of the relationship between fertility and labor force participation. This leads to the conclusion that other aspects of labor force participation, beyond number of hours worked, must play a role in the decision-making process. Furthermore, although the opportunity costs associated with having a child are important to evaluate how a woman makes a fertility decision, it does not give the whole picture. This article shows that preferences are a very significant aspect of fertility behavior. This implies that a woman's preferences for both labor force participation and fertility are influencing her decisions regarding whether or not to have a child.

The previous articles have verified the validity of the premise of my research, but the theoretical approach taken in this paper has remained untouched. The following articles utilize Bayesian statistics to model various aspects of fertility. This usage of Bayesian statistics creates support for and reinforces the reasoning behind my decision to use this approach.

5.4 "Probabilistic Projections of the Total Fertility Rate for All Countries" (Alkema, Raferty, etc., 2011).

In "Probabilistic Projections of the Total Fertility Rate for All Countries", the authors use a Bayesian model to create a forecast of the total fertility rate by nation. Through this model they planned to create a better method for the United Nations to predict fertility rates and, implicitly, population. Their rationale for choosing a Bayesian model was that the fertility pattern found across nations lent itself to a model in which parameters could be treated probabilistically. Furthermore, in much the same way as this research, the total fertility rate projections can, and should, be updated periodically which is easily done through Bayes' Theorem. Overall, the researchers chose this model because it allows for assessment of uncertainty in the estimation of various parameters.

The authors began by identifying a noticeable pattern in fertility rates in each country. From examining past data, they noticed that all countries began at a relatively high and stable fertility rate then entered a period where they transitioned from the high rate to either replacement level fertility or below. The final phase consists of the time post-transition, when fertility rates fluctuate around replacement level. In order to overcome the different behaviors of the fertility rate which characterize each phase, the authors created three different models to predict total fertility rate depending on which phase the country in question is in. Each individual model is dependent on five country-specific parameters that were determined based on the United Nation's current methodology for predicting fertility rates. As previously mentioned, each country specific parameter is treated probabilistically. Furthermore, as more information about these parameters is gathered, the model can be updated according to Bayes' Theorem.

To test their model, the authors chose six countries; Italy, China, the United States, India, Israel, and Mozambique. The group of countries included two countries currently below replacement level fertility and thus in the third phase (Italy and China), one that had recovered from low fertility and is currently fluctuating around replacement level (the United States), one experiencing rapid decline in phase two

(India), and two countries which have experienced very slow decline in fertility levels in phase two (Israel and Mozambique). This group was chosen so that the model could be tested in countries that have not only experienced different fertility rate behavior, but also have very different country specific parameters, such as economic behavior, infant mortality rate, and quality of life. Using data from these countries and the UN projections for total fertility rate, the researchers estimated their model using fertility rates from 1980-2010.

Table 5.4.1 below shows the percentage of UN estimates that fall outside the 95% projection interval. If the Bayesian prediction model was performing well, we could expect approximately 2.5% of the values to fall above the 95% interval and approximately 2.5% to fall below. Despite the fact that the model predicts UN fertility rates well, due to the fact that the projection intervals are based on the United Nation's estimates of the total fertility rates, it is more reasonable to assume that the model creates a projection of future estimates of the fertility rate. There are possibilities for the model to be improved so that it could serve as a better tool to estimate fertility, however, much of the information that would be needed to improve the model is too difficult to reliably obtain. Nevertheless, the model's ability to evaluate the estimates created by the United Nations could serve as a valuable tool to modify any estimates that are determined to be unlikely by the model.

Table 5.4.1

	Proportion of Observations Above 95% Projection Interval	Proportion of Observations Below 95% Projection Interval
1980-1985	0.05	0.01
1985-1990	0.03	0.05
1990-1995	0.04	0.07
1995-2000	0.03	0.10
2000-2005	0.02	0.07
2005-2010	0.02	0.04

Although the article has overarching themes of fertility and Bayesian methods that are similar to the concentration of this research, the article focuses on fertility on a nationwide level while this research addresses the fertility of each individual woman. However, despite this inherent difference, the methodology presented in the article has important implications for this more individualized research. Primarily, the motivations behind the authors' choice of Bayesian methods in the article were similar to the motivations for this research. Fertility is inherently something that changes based on other circumstances. On a nationwide level, the authors point out that certain country-specific parameters have an effect on the total fertility rate. In a similar manner, woman-specific parameters have an effect on fertility decisions. Furthermore, it is important to both models that the parameters be treated probabilistically. In much the same way that a country's economy experiences some degree of randomness, a woman's income, for example, experiences randomness as well. Overall, the methodological approach presented in this article, reinforces the applicability of Bayesian methods to predict fertility.

5.5 "Women's Employment and Fertility: A Welfare Regime Paradox" (Hilgeman & Butts, 2009)

Another article that uses Bayesian methods to address issues related to fertility is "Women's Employment and Fertility: A Welfare Regime Paradox" (Hilgeman & Butts, 2009). Although, this article focuses on a more macroeconomic issues than my research, it also includes microeconomic analysis at an individual level. Futhermore, Hilgeman and Butts examine the relationship between labor force participation and fertility in certain European countries using a Bayesian model.

Beginning in 2000, researchers began to notice an interesting fertility trend occurring in Europe. It seemed that countries in Southern and Eastern Europe, which are typically associated with being more familialistic, were experiencing lower fertility. In their article "Women's Employment and Fertility: A Welfare Regime Paradox", Christin Hilgeman and Carter T. Butts seek to address this paradox (2009). In order to fully address this topic, they create a series of hypotheses they intend to investigate. The first hypothesis is that part time work will affect fertility in a similar way as not participating in the labor force. Secondly, they hypothesize that countries with high total female labor force participation will see more women in the labor force having children because it is viewed as more socially acceptable. They also believe that family leave will have no effect on fertility and enrollment in childcare will be positively associated with fertility rates. To test the relationship between family leave, childcare enrollment, female labor force participation, and fertility, they create a Bayesian model.

They formulate a model using the realized fertility as the dependent variable and individual-level variables, including marital status, education, employment, parental co-residence, and a variable to determine attitude, as well as macro-level variables, including percentage of children in childcare, family leave, and total female labor force participation. The model is formulated as a hierarchical Bayesian model in order to account for both individual- and country-level variables. They begin with a function to account for differences in national fertility rates that is determined by the country in which they live. This is then amended according to certain age and individual effects.

Hilgeman and Butts take a step-by-step approach to develop the model. Firstly, they address the likelihood of the country-specific fertility hazards, β_{0i}, where *i* corresponds to the country. It is assumed that this is normally distributed with a mean determined by a linear combination of country-level covariates and coefficients, represented by Z_i, γ is a parameter vector, and σ_γ^2 is a non-negative real parameter. This likelihood is shown below,

$$p(\beta_{0i}) = N(\beta_{0i} \mid Z_{i\bullet}\gamma, \sigma_\gamma^2).$$

Next, they examine the likelihood of country-specific covariate effects, β_{ij}, associated with individual-level covariate set, *X*. Again assuming normal distribution, this likelihood is given below.

$$p(\beta_{ij}) = N(\beta_{ij} \mid \mu_i, \sigma_i^2).$$

In this equation, μ_i and σ_i^2 are the mean and variance for the population from which the country-specific effect, β_{ij}, is drawn. Lastly, the likelihood of the realized fertility vector, *y*, given the country-specifc and individual-level covariates, age-specific hazard parameters (α), and amount of time spent in an age interval (ε) for individual *i*, is assumed to be a Poisson distribution and is given by the following equation,

$$p(y_i \mid \alpha, \beta, \varepsilon, X) = Pois(y_i \mid \exp(\beta_{0c_i} + X_{i\bullet}\beta_{\bullet c_i})(\varepsilon_i^T \alpha)).$$

Finally, based on all of these prior distributions, the posterior distribution can be represented as follows according to Bayes' Theorem,

$$p(\alpha, \beta, \mu, \sigma, \gamma, \sigma\gamma \mid y, \varepsilon, X, Z) \propto \left(\prod_{i=1}^{n} p(y_i \mid \alpha, \beta, \varepsilon, X)\right)\left(\prod_{i=1}^{n_b}\prod_{j=1}^{m} N(\beta_{ij} \mid \mu_i, \sigma_i^2)\right)\left(\prod_{i=1}^{m} N(\beta_{0i} \mid Z_{i\bullet}\gamma, \sigma_\gamma^2)\right)\sigma_\gamma^{-2}$$

The method the authors employ in this article goes more in depth than the one discussed in this research due to the hierarchical approach. However, the basic methodology behind both approaches is the same. This research examines the likelihood of having a child given labor force opportunity costs while the authors of the article examine the likelihood of realized fertility given country- and individual-specific data. In this article, the authors analyze various country- and individual-level factors that they believe have an effect on fertility. This research, on the other hand, looks more specifically at a woman's individual decision-making process. So, instead of focusing on exogenous factors that affect fertility, this research attempts to address those factors that are endogenous to the model that the woman takes into consideration when deciding whether or not to have a child. However, the motivations behind using Bayes' Theorem are very similar. The Bayesian approach allows for the distribution to be updated as more information is gathered, and the initial prior distribution can be formulated with very little information.

Using the World Values Survey from 1995-1997 and the European Values Study from 1999-2000, Hilgeman and Butts gathered individual-level data on marital status, employment, education, and parental co-residence for the United States and various countries within Europe. The OECD provided county-level data. The authors then ran hypothesis tests on their models to determine 95% posterior probability intervals. In all of the countries addressed, fertility is most associated with marital and employment status. However, as evidence against their first hypothesis, both part time and full time employment have a negative effect on fertility, while staying out of the labor force has a positive effect. Similarly, the results do not support their second hypothesis. Female labor force participation has a negative moderating effect on fertility. Finally, the effect of family leave on fertility is not significant in accordance with their third hypothesis. These results are presented in table 5.5.1 below.

Table 5.5.1

	Point estimate mean	95% Probability Interval	
		Lower	Upper
Individual-level effects			
Full-time employment	-0.36	-0.45	-0.28
Part-time employment	-0.16	-0.23	-0.09
Unemployed	-0.10	-0.22	0.02
Country-level effects			
Family leave	-0.004		

Although the model presented in this article uses both macro- and micro-level variables, the intuition behind its use is similar to the intuition for my research. Hilgeman and Butts use Bayesian statistics so that they can update the distribution as more information is gathered about fertility. In my research the concept is that as a woman gains information about the opportunity costs of having a child, the likelihood that she has the child will change. Overall, this article validates the use of Bayesian statistics and reinforces previously determined relationships between female labor force participation and fertility.

5.6 Conclusion

The literature reviewed in this chapter serves as background and support for the theory presented in the paper as a whole. Throughout the years, research concerning fertility and labor force participation has evolved. The field began with Gary Becker's work, in which fertility was first approached with an economic viewpoint, and has progressed past research uncovering the interconnectedness of labor force participation and fertility decisions. Now, much of the research is focused on a more in depth analysis of the many aspects of fertility in regards to labor force participation. My research is a proposition of a way in which the research about fertility decision-making can be expanded based on commonly accepted economic models and newly tested Bayesian statistical methods. The literature presented here motivated my decision to investigate new possibilities for this concentration and has inspired ways to further my own research that are presented in the following chapter.

6. Implementation and Further Research

Given the theory developed out in this paper, there are many possibilities for implementation and further research. One possible extension of the research would be to use Bayesian methods to create an analysis of the economic concept of utility in terms of the most likely outcome of the fertility decision-making process Another possibility to expand the theoretical approach would be to create hypothetical functions for the production possibility frontier and utility function over time. Lastly, doing a regression analysis to test the applicability of the Bayesian model could expand the research. The following sections elaborate on the potential for further research.

6.1 Decision Analysis

Another possible extension of the model would be to use Bayesian decision analysis to examine the interaction with the Bayesian model and the utility function. In the chapter on Bayesian statistics, the process of Bayesian decision analysis was briefly explained. This process determines the expected value for a parameter, θ, given the utility functions of the consequences of various possible actions. For the purposes of this research, each consequence q_i would be a combination of amount of hours worked and the number of children a woman has. Thus, as previously discussed, a woman's utility will be a function of children and work, represented by U(c, w). So, the bundle of c_1 children and w_1 hours worked will be represented by consequence q_1. Therefore, similar to the economic approach to utility functions, if a particular value for the utility is chosen and the bundle (c_1, w_1) and the bundle (c_2, w_2) both yield that utility, then the person is indifferent between the two consequences and both bundles would fall upon the same indifference curve. Hence, if the bundle (c_1, w_1) is represented by consequence q_1 and the bundle (c_2, w_2) is represented by consequence q_2, then U(q_1) = U(q_2).

Although Bayesian statistics often uses loss function in the place of utility functions to minimize loss, it would be fairly straightforward to use the original utility function. For example, suppose a woman has the very simple utility function $U(c,w) = kcw$ where k is some constant. Further suppose that both c and w are functions themselves. For example, $c = l\theta$ and $w = r(1-\theta)^2$ where both r and l are constants and θ is the probability that a woman will have a child. Then, the utility function can be further expanded to be $U = C\theta(1-\theta)^2$ where C is the product of the constants. Given this utility function, the expected utility would be

$$E(U) = \int_0^1 C\theta(1-\theta)^2 h(\theta \mid x_1, x_2, \ldots, x_n)\, d\theta \, .$$

Furthermore, the expected value of θ could also be estimated using the method described in the chapter on Bayesian statistics.

This method could potentially offer an approach to determining the utility at which the bundle of child services and work would be maximized at the point of tangency with the production possibility frontier. This could mean a move operationalized approach versus using graphs to approach the issue of utility theoretically. This method follows the economic assumption that if various bundles of child services and work yield the same utility, the woman would be indifferent between the two bundles and it involves utilizing the assumption made in this research that both child services and work are a function of θ. Further, it applies the concept of production possibility frontiers in the use of the posterior distribution function and, thus, the opportunity costs associated with having children.

6.2 Creating Functions

Given the assumption that both the utility function and production possibility frontier are dynamic, another possibility for expanding this research would be to create functions for these curves. Using the assumptions laid out in the previous chapter that both child services and work are a function of the parameter θ, creating functions would allow examination of the changes occurring due to differing values of θ over time. The utility function would, therefore, be of the form $U_t(c(\theta), w(\theta))$. So, both child services and work would still be functions of θ and the utility would also change over time. A function for the production possibility frontier would also be dependent on time. Since the production possibility frontier is linear, in this case, the production possibility frontier could be written as $f_t(c(\theta), w(\theta))$ such that f is a linear function over time. This would allow both child services and career to vary over time as well as the slope of the function, as is discussed in chapter 3.

6.3 Testing the Model

Running a regression analysis would allow us to test the applicability of the model in a real world application. Data would have to be gathered about various women's preferences towards having children and participating in the labor force and the opportunity costs they might incur if they were to have children. It would be possible to create an estimate of the opportunity costs by collecting data on potential promotion opportunities and income and then estimating the possible loss in future income if a woman were to have a child and the potential cost of human capital depletion. Futhermore, data would have to be gathered on marital status, how many children are already present in the household, and whether or not a woman had a child after the previous data collection. This data would have to be collected at various, pre-determined time intervals in order to allow the Bayesian model to be updated multiple times throughout a the women's fertile years.

The data would be entered into the model at each time interval to predict the likelihood of a particular woman having a child at that given time. At each subsequent time, the process would be repeated. Then, the distributions yielded by the model could be compared with the women's actual fertility behavior. In theory, a distribution for θ that places a higher probability on θ being close to one, should be associated with a woman having a child. Conversely, if the distribution reveals that the woman is unlikely to have a child we would expect that the woman would not have had a child between the time periods.

A regression analysis could be run to compare what actually happened with what the model predicted as the likelihood for having a child. Using this comparison, we would get a clear picture of the ability of the Bayesian model to predict the likelihood that a woman has a child. We would also gain information about the degree of accuracy of the model and it would allow us to investigate potential improvements. Although the data necessary would be more in depth than most data collected, Shreffler and Johnson found multistage data collected by the National Survey of Family and Households and Kahn and Whittington used data collected individual women from the Puerto Rico Fertility and Family Planning Assessment (Shreffler & Johnson, 2013) (Kahn & Whittington, 1994). Hence, similar data is available and could be used to test this model.

6.4 Conclusion

These potential extensions could be used to enhance the applicability of the model as well as to test how applicable it really is. Using these expansions would allow a greater understanding of the many ways in which the variables within the fertility decision-making process interact and affect each other. Further, as the model is expanded and tested, improvements could be made to enhance the model and make it more representative of the actual implicit process women go through when making decisions regarding their fertility in relationship to their labor force participation.

7. Conclusion

The model created in this paper is a formalization of the decision-making process a woman in the labor force uses when deciding to have a child. Utilizing production possibility frontiers and indifference curves, it is based on models that have been widely accepted in feminist and classical economics for many years. To account for the more dynamic nature of the issue that is often ignored in other literature on the subject, a Bayesian model was used. Overall, this fertility decision-making model is a new angle on the well-studied topic of fertility and labor force participation.

Similarly to many of the previously suggested models, this particular model has its limitations. One of these limitations is common in many economic models. The model presented in this paper only considers two possible activities in the production possibility frontier and indifference curves: having and raising children and participating in the labor force. Obviously, women have many other options. The use of a dualistic model was for simplicity's sake. Further, this also means that the model only applies for women in the labor force. The only opportunity costs being considered are those that relate to labor force participation. However, it would be possible to expand the model to allow for other opportunity costs or variable to be included. This would require determining distributions for each added variable and would affect the conjugate prior as well as many other aspects of the model. Another limitation is that the model assumes that women are considering their labor force participation in their decision to have children. It is possible that some women do not consider the possible effects that a child may have on their career. This may be mitigated by the inclusion of indifference curves because many of the women who are not considering labor force opportunity costs probably place higher value on having children. As any model does, the model presented in this research makes many necessary assumptions.

Taken as a whole, the model is a simplification of a complicated and emotional issue. However, this model is an improvement on older models in many ways. Firstly, as previously mentioned, the vast majority of older models examine fertility decisions in a static framework. In my opinion, that is an extreme oversimplification. The model created in this paper includes the changing nature of both preferences and opportunity costs. The use of Bayes' Theorem allows us to take into account the fact that women base their decisions on prior decisions as well as new information. Further, the definition of the parameter θ as the likelihood of having children instead of the act of having children includes the uncertainty at play in fertility. Many models make the assumption that a woman who decides to have a child will be immediately successful at conceiving said child. Despite its limitations, in my opinion this model is an improvement upon the models presented in previous research. It approaches a topic that many others have researched and modeled from a new angle.

Gaining a deeper understanding of the role that labor force participation opportunity costs play in women's fertility decisions could help to shape future public policy decisions. For example, the model proposes that women consider their foregone income when deciding to have children. This could potentially lead to many well-educated, well-paid women deciding to delay or even forgo having children. However, introducing guaranteed paid maternity leave could negate the importance of forgone income. Overall, the model presented in this paper expands the understanding of fertility decisions. Using the information gathered from this model to create public policy could lead to more equality for women participating in the labor force.

References

Alkema, L., Raferty, A. E., Gerland, P., Clark, S. J., Pelletier, F., Buettner, T. & Heilig, G. K. (2011) “Probabilistic projections of the total fertility rate for all countries”. *Demography*. Population Association of America.

Arroyo, C. R. (1993). *Economic approaches to modeling fertility determinants.* The World Bank.

Babu, G. J. (2012). “Bayesian and frequentist approaches”. Online Proceedings of the *Astronomical Data Analysis Conference (ADA VII).*

Becker, G. S. & Lewis, H. G. (1974). “Interaction between quantity and quality of children”. *Economics of the family: Marriage, children, and human capital,* 81-90.

Berger, D. O. (2001). *Bayesian and conitional frequentist hypothesis testing and model selection*. PowerPoint. VII C.L.A.P.E.M. La Habana, Cuba. < http://www.stat.duke.edu/~berger/talks/cuba-preve.pdf>

Blau, D. M. & Robins, P. K. (1989). “Fertility, employment, and child-care costs”. *Demography, 26*(2), 287-299.

Cowles M. K. (2013) *Applied Bayesian statistics.* New York, NY: Springer Science and Business Media.

Cramer, J. C. (1979). “Employment trends of young mothers and the opportunity cost of babies in the United States”. *Demography 16*(2), 177-197.

De Tray, D. N. (1974). “Child quality and the demand for children”. *Economics of the family: Marriage, children, and human capital,* 91-119.

Gustafsson, S., & Wetzels, C.M.M.P. (1999). *Squeezing birth into working life. household panel data analyes comparing Germany, Great Britain, the Netherlands, and Sweden*. 93-110.

Hilgeman, C. & Butts, C. T. (2009). “Women’s employment and fertility: a welfare regime paradox”. *Social Science Research 38*, 103-117.

Hondroyiannis, G. (2010). “Fertility determinants and economic uncertainty: An assessment using European panel data”. *Journal of family economic issues, 31,* 33-50.

Kahn, J. R. & Whittington, L. A. (1994). "The transition to parenthood in Puerto Rico: Occupational status and the timing of first births". *Population research and policy review, 13*, 121-140.

de Laat, J. & Sevilla-Sanz, A. (2011). "The fertility and women's labor force participation puzzle in OECD countries: The role of men's home production". *Feminist economics, 17*(2), 87-119.

Lee, P. M. (2012) *Bayesian statistics: An introduction.* (4 ed). Somerset, NJ: Wiley.

Lehrer, E. & Nerloce, M. (1986). "Female labor force behavior and fertility in the United States". *Annual Review of Sociology, 12*, 181-204.

Morgan, B. W. (1968). *An introduction to Bayesian statistical decision processes.* Englewood Cliffs, NJ: Prentice-Hall, Inc.

Press, S. J. (1989) *Bayesian statistics: Principles, models, and applications.* New York, NY: Wiley.

Schmertmann, C. P., Cavenaghi, S. M., Assuncao, R. M. & Potter, J. E. (2013). "Bayes plus Brass: Estimating total fertility for many small areas from sparse census data". *Population studies: A journal of demography.*

Schultz, T. P. (1969). "An economic model of family planning and fertility". *Journal of political economy, 77*(2), 153-180.

Shreffler, K. M. & Johnson, D. R. (2013). "Fertility intentions, career considerations and subsequent births: The moderating effects of women's work hours". *Journal of family economic issues, 34,* 285-295.

Tullock, G. & R. B. McKenzie. (1985). "Child production". *The new world of economics, 4*. Homewoord, IL: Richard D. Irwin Inc.

Waite, L. J. & Stolzenberg, R. M. (1976). "Intended Childbearing and Labor Force Participation of Young Women: Insights from Nonrecursive Models". *American Sociological Review 41*(2), 235-252.

Willis, R. J. (1973). "A new approach to the economic theory of fertility behavior". *Journal of political economy 81*(2), S14-S64.