2013-06-20

Modelling Infertility with Markov Chains

Rebecca Dorff
Brigham Young University - Provo

Modelling Infertility with Markov Chains

Rebecca Dorff

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Jeff Humpherys, Chair
Chris Grant
Kevin Seppi

Department of Mathematics

Brigham Young University

June 2013

ABSTRACT

Modelling Infertility with Markov Chains

Rebecca Dorff
Department of Mathematics, BYU
Master of Science

Infertility affects approximately 15% of couples. Testing and interventions are costly, in time, money, and emotional energy. This paper will discuss using Markov decision and multi-armed bandit processes to identify a systematic approach of interventions that will lead to the desired baby while minimizing costs.

Keywords: Infertility, Medical Diagnosis, Markov Decision Processes, Multi-armed Bandits

Acknowledgments

I'd like to first thank my adviser Dr. Humpherys for working with me, despite his incredibly busy schedule. He helped me open more doors than I can count and gave me opportunities that I never would have dreamed of. I'd also like to acknowledge my committee members, Dr. Chris Grant and Dr. Kevin Seppi from Brigham Young University for their time and comments on this work. Joseph B. Standford, MD, from the University of Utah came up with this fascinating problem, for which I am forever grateful. It was the perfect thesis for me. I couldn't have done it without the Markov Seminar participants, James Taylor, Jeff Hendricks, Abe Frandsen, and especially Mason Victors. I'd like to give a shout out to those at the BYU math department who significantly impacted my life, Dr. Rod Forcade, Dr. Pace Nielson, and Dr. Gary Lawlor. Finally, I'd like to thank my parents for their wonderful support.

Contents

List of Tables

List of Figures

Chapter 1. Introduction

One of the most prevalent diseases in the world is infertility. After six months of attempting to conceive, about 20% of couples are considered at least slightly infertile [28]. According to a 2010 analysis of 277 health surveys of women between the ages of twenty and forty-four, in developing countries, this number increases to 25% [37]. Only half of them were able to conceive within five years, leaving 12% of women unable to have a child. While these numbers are on the high end, even the conservative estimates are still considered quite high, with 10% of women considered infertile within a year and 5% of women considered sterile after four years [28, 31]. The wide range may be attributed to the fact that only about 50% [31] of infertile couples seek medical care. In the United Kingdom, one in every seven woman, or 14%, seek medical treatment for infertility [55].

Not only does infertility affect a large number of women, the number is increasing. A woman's ability to become pregnant diminishes with age. As more women wait until they are older to have children, their fecundity, or ability to conceive, decreases. Because of this, infertility has increased by four percent since 1980 [2]. Another study showed that fewer women are trying to have children, indicating a reduction in primary and secondary infertility, but since the population has increased, the number of couples affected by infertility has increased by 16% since 1990 [37]. To put this in perspective, about one in every twelve people, or 8.3% of the United States, suffers from diabetes [38]. Heart disease, the leading cause of deaths worldwide, causes one out of every four deaths [24].

Despite affecting so many people, infertility is not widely discussed, mostly because it is such a sensitive subject. In some societies, women who are unable to bear children are considered evil or cursed [34]. In a study initiated by Schering-Plough and Merck & Co. in 2009 of 585 men and women, sixty-one percent of couples hid their infertility from family and friends. Almost half of the women didn't tell their mothers. Around 61% of infertile couples experience rejection or isolation. The couples also face stress, anxiety, and depression, among others. In 2011, Redbook, a magazine for women, launched a video called, "The Truth of

Trying", to raise awareness of infertility and help people become more open [19].

The complexity of the problem makes it difficult to diagnose and testing can only reveal certain problems. Up to 30% of infertility is unexplained [53, 29, 19]. There is no set procedure for determining the causes of infertility. It is relatively easy to test hormone levels and semen counts, but after that, tests can become expensive and invasive. Since diagnosing and solving infertility is such a large problem, it can be difficult for doctors to know how to approach the problem. While talking with couples who've struggled with infertility, we've noticed that sometimes the doctors assign treatments based solely on how much the couple can spend or what they feel will work.

The purpose of this paper is to create an objective approach to treating infertility by modeling infertility and simulating treatments. Modeling and simulating medical treatments is gaining popularity. In "Modeling and Simulation in the Medical and Health Sciences," Sokolowski and Banks wrote "there is now, more than ever, a need for partnership between engineer and health care professionals; this is if engineers as developers and medical professionals as users are to meet the six goals outlined by the Institute of Medicine" [54]. By modeling and simulating, we hope to achieve better results than costly and time intensive clinical trials.

The goal of our program is to to maximize the chance of a couple getting pregnant in a given time frame, while staying within their budget. To do this we use a multi-armed bandit approach. Multi-armed bandits have gained popularity in simulating clinical trials due to ethical issues, monetary and time constraints, and reduced avoidable treatment failures [49]. The difficulty of using a straight multi-armed bandit problem with infertility interventions is that cost is a serious issue. We can't simply choose the treatment that will always yield the highest result. For example. IVF might increase the chance of getting pregnant for most couples, but at a base rate of $13000, not many couples can afford to try it several times. To compensate for this, we also incorporate a dynamic programming model similar to the Cake Eating problem to account for the budget a couple has, as well as how much time they

have. In this problem, time becomes a currency since the ability of a woman to become pregnant decreases with age. A thirty-eight-year-old woman cannot afford to spend a year trying naturally to get pregnant, whereas a twenty-eight-year-old woman can. We combine the the bandit solution for the entire population with the bandit solution for the individual and the amount of time and money the couple has to create an index that determines which intervention should be tried at a given time period.

The rest of the paper will be as follows:

Chapter 2 discusses the medical background of infertility, including the causes and treatments of infertility. Chapter 3 introduces the mathematical background needed to understand the model. Chapter 4 describes the model. Chapter 5 discusses the results and analysis the model. Chapter 6 concludes the paper.

Chapter 2. Medical Background

Infertility is a complex problem. There are many different causes and almost as many definitions. In this chapter, we will discuss the terms and definitions associated with infertility. We will then explain the reproductive cycle. The next section identifies the places where infertility occurs in the reproductive cycle

2.1 Definitions and Terms

Fertility, the capacity of producing children, is measured by the number of offspring. It is often confused with fecundity, which is the potential reproductive capacity of a being. Infertility has several definitions., including a demographic definition and an epidemiology definition. The clinical definition, as given by WHO, the World Health Organization, is "a disease of the reproductive system defined by the failure to achieve a clinical pregnancy after 12 months or more of regular unprotected sexual intercourse " [46]. This definition was expanded by Gnoth et al. to include that intercourse must occur during the "fertile phase of the menstrual cycle" [28]. The epidemiology definition increases the amount of time until infertility to be two years. It is this definition that is used in epidemiological studies while the clinical one year definition is used in practice. There is reluctance to accept this definition of infertility as a disease from governments and private health organizations [32]. This, along with the popular belief that infertility indicates sterility, has led to a new term, subfertility, which is often used interchangeably with infertility. Subfertility refers to reduced fertility over a period of time. This term is widely used in practice, but a specific time frame hasn't been formally defined for the different grades of subfertility. Both Habbema et al. and Gnoth et al. propose formal definitions for infertility and subfertility, although the WHO definition given above seems to be the prevailing one in the literature [32, 28].

Even though infertility does not imply sterility, the inability to ever conceive, the general population tends to believe that infertility leads to sterility [32]. One cause for the belief

that infertility is equivalent to sterility is the WHO definition of primary infertility, which is when a woman is unable to ever have a live birth. Primary infertility is sometimes used interchangeably with infertility. In the literature, primary infertility typically means that the woman has had no prior child, regardless of whether or not she is capable of having future children [46, 41]. Secondary infertility refers to couples who have been pregnant at least once, but are unable to conceive a second time [46]. Again, like primary infertility, it is often used to indicate that the woman hasn't had a second birth, not whether or not she is capable [41, 32].

There is some debate about whether infertility is a disability. According to WHO, infertility causes an impairment of function, making it a disability. Infertility in women was ranked the 5th highest serious global disability for people under the age of sixty. Under this definition, "health care falls under the Convention on the Rights of Persons with Disability" [46]. Despite this, infertility is not legally considered a disability even though treatments and psychological effects of infertility allow infertile people to use the Family and Medical Leave Act (FMLA) [34]. Usually treatments and interventions are not covered by insurance, leaving a huge financial burden on the couple.

2.2 Female Reproductive Cycle

The menstrual cycle in women lasts approximately twenty-eight days. It is broken up into two phases. The Follicular Phase starts the first day of the last menstruation cycle and lasts anywhere from 7-40 days. The Luteal Phase starts the first day of ovulation and lasts 14-16 days.

At the beginning of the Follicular Phase, low estrogen levels prompt the hypothalamus, the part of the brain that links the nervous system to the endocrine system, to send a signal to the pituitary gland to release the follicle-stimulating hormone, known as FSH. FSH prompts several follicles, which are fluid-filled cysts in the ovaries, to develop into mature eggs. As the follicles mature, they release estrogen. The high levels of estrogen tell the

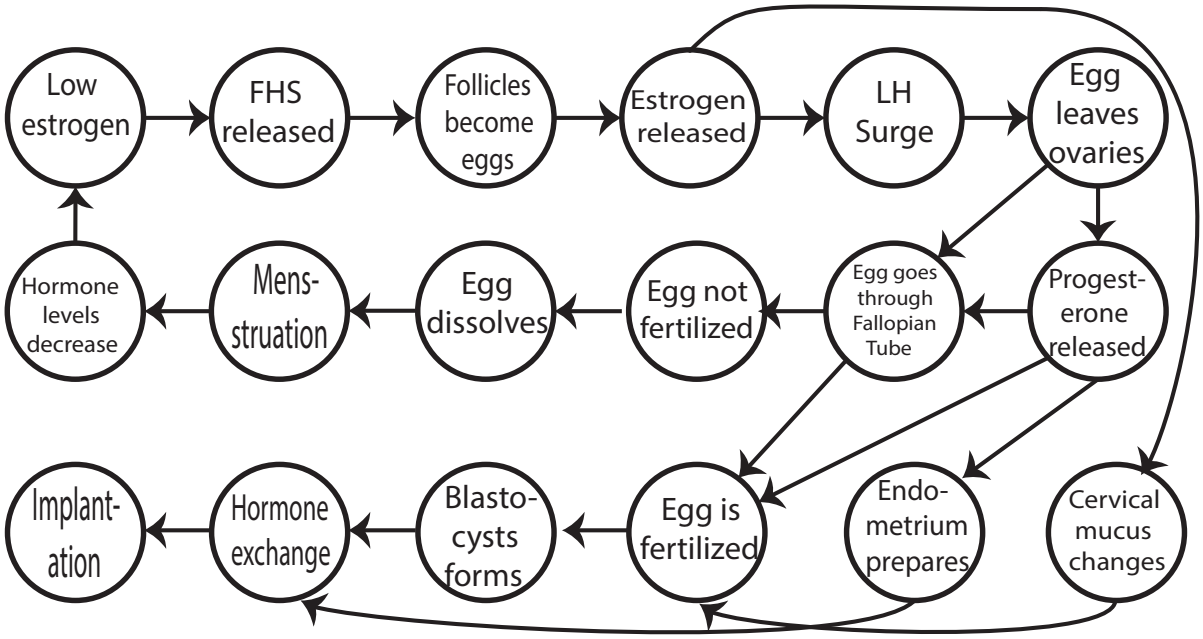

Figure 2.1: The Menstrual Cycle

hypothalamus and pituitary gland that mature eggs, or ova, are ready. This causes LH, the luteinzing hormone, to be released in what is known as the LH surge. The LH surge causes the dominating egg, or in some cases, eggs, to push through the ovary wall into the Fallopian tube. This is typically known as ovulation and is the beginning of the Luteal Phase. In the rare event that two eggs are dominating, fraternal twins may result.

During the Luteal Phase, the coprus luteum, the follicle that released the dominating egg, starts releasing the hormone progesterone. Using progesterone, the endometrium, or uterus wall, starts storing blood and nutrients for the egg. Once in the Fallopian tube, the egg has to be fertilized within 24 hours, or it disintegrates. Sperm can live up to 2 days in the Fallopian tubes, depending on their quality and the state of the cervical mucus. In order to reach the Fallopian tubes, which are stocked with nutrients, the sperm must travel through the cervix and cervical mucus. High levels of estrogen change the consistency of the mucus, as well as increase the pH, so that sperm can survive. Progesterone helps sperm penetrate the egg by shedding the layers of protein surrounding the sperm.

If the egg is fertilized, it will form a blastocyst, a hollow shell of cells that develop from

within. The outer wall will become the placenta. The zygote, or fertilized egg, continues down the Fallopian tube toward the uterus, about a four-day journey. Once in the uterus, the blastocyst contacts the endometrium. They exchange hormones, which allow the blastocyst to connect to the uterine wall, a process called implantation. During implantation, the cervix is closed by a mucus plug. If the egg is not fertilized within 24 hours, it dissolves into the uterus. After 12-16 days, the endometrium lining, the blood and nutrients collected during the Luteal Phase, breaks down and is released during menstruation. Hormone levels decrease and the cycle begins again [9].

2.3 Infertility in Reproductive Cycle

There are five main steps in the reproductive cycle where infertility can occur [41].

(i) The ovaries don't produce eggs.

(ii) The eggs don't travel from the ovaries to the uterus.

(iii) The egg doesn't get fertilized.

(iv) The zygote doesn't attach to the endometrium.

(v) The zygote doesn't survive after implantation.

Each of these problem spots and their possible causes will be discusses in further detail.

Ovulation Problems. Ovulation problems arise when the ovaries don't release eggs regularly. When no eggs are released in a minimum three month period, we define the problem as anovulation. Anovulation basically means that ovulation doesn't take place, so it is often characterized by irregular or missing menstrual periods, known as amenorrhea. Anovulation can be caused by several factors: hormone imbalances, eating disorders, poor nutrition, alcohol, stress, smoking, excessive exercising, obesity, diabetes, cancer, tumors, thyroid disease, Premature Ovarian Failure (POF),and cysts, such as Polycystic Ovary Syndrome (PCOS) and endometriomas (ovarian cysts) [6].

We consider each of these problems a little further. A hormonal imbalance is considered the biggest cause of anovulation, accounting for 30% of female infertility. Because each step of ovulation is initiated by a hormone release, a lack or excess of estrogen, FSH, or LH can affect the process. Too much testosterone can also lead to anovulation. When even one hormone is out of balance, the other hormones also become out of balance [57]. Hormone imbalance often occurs when the hypothalamus or pituitary gland malfunctions. The thyroid diseases hypothyroidism (underactivity) and hyperthyroidism (overactivity), create hormone imbalances from the thyroid hormones interacting with the repoductive system, which leads to irregular periods. Hypothyroidism can also lead to weight gain, high levels of prolactin, and ovarian cysts. Between 2 and 4% of women suffer from thyroid diseases [35].

Lifestyle factors can also have an impact on ovulation by affecting these hormones. Estrogen is produced in fat cells, so having excessive fat contributes extra estrogen, throwing hormones out of balance. On the other hand, having too little fat results in not enough estrogen being released. Thus obesity, eating disorders, and excessive exercise all affect ovulation by creating hormone imbalances. Alcohol, even in small amounts, can causes anovulation by affecting the liver, which regulates the amount of estrogen in the body, leading to excessive estrogen. Alcohol can also cause amenorrhea, the lack of a menstruation cycle; Luteal Phase Defect, discussed later; and hyperprolactinemia, in which a woman has high levels of prolactin, a hormone produced by breast feeding [33]. Smoking increases the amount of FHS, accelerating follicular depletion, and decreases the amount of estrogen [43].

PCOS, or polycystic ovary syndrome, affects between $5-10\%$ of women and is thought to account for large amount of anovulation. PCOS refers to a hormone imbalance, specifically estrogen, progesterone, and tetosterone, that lead to eggs not maturing. Women with PCOS often develop insulin resistance, which also contributes to irregular periods. The irregular periods, along with immature eggs, cause ovarian cysts to develop [20, 42]. Diabetes, or any insulin resistance, affects insulin, causing a hormone imbalance, which then triggers an imbalance in estrogen, progesterone, and testosterone. Insulin resistance also can contribute

to miscarriages [38]. Premature Ovarian Failure (POF), or primary ovarian insufficiency, affects 1% of women and is the loss of function in the ovaries [14]. It results in abnormal amounts of estrogen and eggs that are not released. It is possible to have a spontaneous pregnancy with POF since irregular or occasional periods may occur.

Blockage. The most common reason the egg can't reach the uterus is that there is something blocking the way. Blocking can result from growths, either fibroids or polyps, in the cervix or uterus, birth defects in the woman, Pelvic Infections (PID), medications, endometriosis, scar tissue, and surgery [6].

Growths in the cervix, uterus, and Fallopian tube block the egg from traveling to the uterus. Uterine fibroids, which are noncancerous tumors, affect implantation when they block the Fallopian tube or are located in the uterine, though this rare [56]. Polyps, abnormal growths from mucus membranes, can occur in the uterus and cervix. Endometrial polyps, polpys in the uterus, occur in about 10% of women. It is conjectured that hormones contribute to their growth. They also can increase the chance of a miscarriage with IVF [13]. Growths and anatomical distortions can occur because of endometriosis. Endometriosis is a condition in which endometrium cells grow outside of the uterus. These cells have no way of leaving the body during menstruation, causing pain, scar tissue, and abnormal growths. Endometriosis is the leading cause of infertility [3].

Pelvic Infections, or pelvic inflammatory disease (PID), is an inflammation of the uterus, Fallopian tube, and/or ovaries which leads to scar tissue. These infections can be viral, parasitic, fungal, and most commonly, bacterial. Scar tissue, caused by PID, endometriosis, or surgery can effectively block the Fallopian tubes, making it nearly impossible for the egg to reach the uterus [6].

No Fertilization. This problem, usually caused by a lack of sperm in the Fallopian tube or the inability of the sperm to penetrate the egg, can be split into two categories, female and male problems [41, 6].

The two biggest male factors associated with infertility are Oligozoospermia and Azoosper-

mia. Both deal with low sperm counts. Oligozoospermia, or oligospermia, refers to semen with a low concentration of sperm while azoospermia indicates no measurable level of sperm in the semen. WHO, the World Health Orginization, defined oligozoospermia to be be less than 15 million sperm/ml [53]. With this threshold, 5% of men have oligozoospermia [17]. Azoospermia affects about 1% of all males [44].

Other sperm disorders include immature sperm and sperm that are unable to move properly. Causes of sperm disorders include infectious diseases; varicocele, which is a swelling of the veins that drains testicles; endocrine disorders; hormonal imbalances; ejaculation problems, including retrograde ejaculation; spinal injuries; medications; surgery; antibodies that attack sperm, common after a vasectomy; anatomical abnormalities like tumors or cancer; sperm duct defects; chromosomal defects (Klinefelter's, cystic fibrosis, Kallmann's syndrome, Young's syndrome, and Kartagener syndrome); Celiac disease; smoking; alcohol; weight; and stress [45]. Other causes for male infertiltiy are erectile dysfunction and hypospermia, a low semen volume.

The main female reason for fertilization not occurring is that the sperm doesn't reach the egg. This usually involves the cervical mucus, which can be insufficient, too thin, too thick, too acidic, or too hostile to sperm. The mucus normally has very low, acidic pH. Estrogen raises the pH of the mucus, which allows the sperm to survive, so low levels of estrogen make it difficult for sperm to survive for longer than a few hours. If the mucus is too thick, the sperm won't be able to move. Thus they won't reach the egg. Medications, including antihistamines even and Clomid can make cervical mucus too hostile. Sometimes, antibodies can occur that attack the sperm. This can be caused by Type 1 Diabetes, immune disorders, and pelvic infections.

No Implantation. In this stage, the zygote, or fertilized egg, fails to implant onto the endometrium. Possible causes include chromosomal issues, which cause the zygote to die; growths and scar tissue, which won't attach to the zygote; the egg doesn't hatch, which means that the shell doesn't soften enough for the egg to implant; infection or a poor

immune system, which kills the zygote; and an unprepared endometrium, which prohibits implantation. This last cause is known as Luteal Phase Defect and generally refers to any abnormal endometrium defects. It is caused by low levels of progesterone or by the endometrium not responding the progesterone.

Miscarriages. When the zygote dies shortly after implantation, it's called a chemical pregnancy. A positive pregnancy test will result because pregnancy tests measure the amount of hCG, human chorionic gonadotropin, a hormone released after implantation, in the blood. Often, the woman will not know that she is pregnant and will experience heavier menstrual bleeding than normal. This is why it is recommended that a woman wait until she's missed her period to take a pregnancy test. It's estimated that between 50 and 70% of all pregnancies end in a chemical pregnancy. Up to 40% of all pregnancies may end in miscarriages because of this [39, 15]. Chemical pregnancies are caused by Luteal Phase Defect, infections, abnormal hormone levels, poor endometrium, blighted ovum (which means that the egg doesn't develop into an embryo and causes about one in two first trimester miscarriages), and uterine growths and abnormalities. The rest are due to chromosomal problems, including poor sperm and egg quality [4, 50].

A miscarriage, or spontaneous abortion (SAB), is the loss of a pregnancy under 20 weeks. $10 - 25\%$ of all clinically recognized pregnancies end in a miscarriage. For healthy women, the chance of a miscarriage is between 15 and 20% [4]. It takes a woman between four and six weeks after a miscarriage to have a period [39]. Most miscarriages, about 75%, occur in the first trimester [15]. Of these, more than 50% of miscarriages are due to chromosomal problems in the fetus [39]. Age also has a large effect the chance of having a miscarriage [4].

Age	%
Under 35	15
35-45	20-35
Over 45	50

Table 2.1: The chance of having a miscarriage based on age

Between 1 and 5% of pregnancies result in a miscarriage in the second trimester, 13-19 weeks [39]. About fifteen percent of these miscarriages are the result of a weakened uterus or cervix, and umbilical cord problems account for 19 percent of second term miscarriages [1, 48].

Since many miscarriages are due to chromosomal issues, 80% of miscarriages are single miscarriages.[15] However, 1% of couples experience repeat miscarriages in a row. Without treatment, $60 - 70\%$ of women with repeat miscarriages will have a successful pregnancy [39]. A woman with a previous miscarriage has a 25% chance of another. The causes for repeat miscarriages are similar to those of single miscarriages, but more women suffer from uterine and placenta problems [39].

Cause	%
Chromosomal	2-4
Uterine or cervical abnormalities	10-15
Antiphospholipid syndrome(blood clots in placenta)	5-10
Hormonal	25-40 (early)

Table 2.2: Percentage of women affected by causes of Repeat Miscarriages

The chance of having a stillbirth, a pregnancy loss after 20 weeks, is 1 in 160 [15, 40].

Cause	% of women affected
Birth defects	15-20
Placenta Problems	25
Poor fetal growth	40
Infections	10-25 (early)
Maternal chronic health	10
Umbilical Cord accidents	2-4

Table 2.3: Percentage of women affected by causes of stillbirths

Other Causes for Infertility. Psychology also plays a role in infertility. One study cited that of women with infertility issues, 40% suffer from poor mental health, 27% from anxiety disorder, and 17% from depression [18]. This is significantly higher than the general average, 3%. Another study showed a strong correlation between depression and infertility. There was a $15 - 54\%$ prevalence of major depression and $8 - 28\%$ with major anxiety in infertile couples [26].

There is some debate on how significant the relationship is and if the psychological issues are a result of the infertility. In a survey of studies on pyschology and infertility, sixteen of the studies showed significant relationship between pre-procedure distress and birth rates while eight of the studies did not [18].

It has been observed that depression is higher in women who are just about to receive IVF, which has been shown through a study by Eugster and Vingerhoets to be more of a psychological burden than a physical one [21]. Several European studies have examined the drop-out rates and reasons for IVF treatments and the biggest cause of dropping out is the emotional burden. It has also been shown that any artificial reproductive treatment has negative short-term emotional impact [26].

2.4 Testing for Infertility

One of the difficulties of infertility is that it can be hard to diagnose. There are several common, easy, and inexpensive test, but then identifying the cause can be difficult. In fact, about 28% of infertile couples are not able to identify the cause [16, 29, 53].

2.4.1 Women. Since there are lots of reasons for irregular ovulation or anovulation, there are several different tests to see if everything is working correctly. The most basic infertility tests check that ovulation is occurring. This can be done at home by measuring morning body temperatures. Progesterone causes body temperature to rise, so there is a small temperature spike around the time ovulation occurs. Also common are home ovulation

tests that measure the amount of LH, luteinizing hormone, in urine. Other hormone tests include measuring the amount of FSH, estradiol (a type of estrogen), progesterone, prolactin, and testosterone. These can be done by analyzing saliva, blood, or urine. Ovarian function tests measure how FSH and estradiol work during ovulation. Luteal Phase testing measures the amount of progesterone and includes more extensive hormone testing.

There are two other simple tests that can be done. A PCT, post coital test, determines if the sperm can move and survive in the cervical mucus. A bacterial screening is often done during the PCT. The other test is an ultrasound. Ultrasounds can measure the thickness of the endometrium, measure follicular growth, check on the ovaries and uterus, and confirm if an egg has been released.

If the tests above, along with the semen analysis, have normal results, further testing can occur. These tests are more invasive and focus on the Fallopian tubes and uterus.

A hysterosalpingogram (HSG) is an x-ray of the uterus and fallopian tubes. Dye is injected into the cervix, which can be used to identify blockages. If the HSG reveals some blockage, a hysteroscopy may be used. During the hysteroscopy, a hysteroscope is put through the cervix into the uterus. This allows the doctor to see any abnormalities, scarring, or growths. A laparoscopy is a minor surgical procedure in the abdomen. A laparoscope is inserted through a cut in the abdomen. It provides a view of the Fallopian tubes, uterus, and ovaries. It is used to identify endometriosis, growths, or scarring, which can be removed with a laser during the procedure. Another test that checks the endometrium tissue is called an endometrial biopsy. It's main purpose is to check the thickness of the endometrium lining to see if implantation can occur [25, 5].

2.4.2 Men. A semen analysis is the most common test for men. It involves taking a sperm sample and checking the count, shape, mobility, appearance, and concentration of the sperm under a microscope. Sperm count tests can be done at home. If the sperm count is low, there will usually be further hormone testing. Urinalysis, a series of urine tests, can be done to locate any infections or sperm in the blood, which is a sign of retrograde

ejaculation. Retrograde ejaculation occurs when the semen travels backwards to the urinary bladder. While the semen analysis will usually catch most male infertility issues, there are a few other tests that are occasionally used. These include an ultrasonography, which checks for blockage in the reproductive tract; testicular biopsy, which tests how the sperm is being produced; hypo-osmotic swelling, which determines if the sperm can penetrate the egg; and sperm agglutination, which tests for sperm clumping [8].

2.5 Treatments

To match the many different causes of infertility, there are many treatments. They fall into the categories of medications, surgical, IUI, artificial reproductive technologies, and pyschological.

Medications. Medications are often given for hormonal imbalances. Listed below are the most common.

Synthetic Human Chorionic Gonadotrophin (hCG), or Gonadotropin-releasing hormone (Gn-RH) analog, is given through the leg to women with irregular ovulation or women who ovulate before the egg is ready. GN-RH affects the pituitary gland to change the timing of ovulation and increase egg maturation.

Synthetic Follicle Stimulating Hormone (FSH) is injected below skin and causes the ovaries to begin ovulation. Synthetic hormones can cause mood swings, ovarian hyperstimulation syndrome, multiple births, and premature deliveries.

Clomidane citrate (Clomid) is an oral medication for women with ovulatory problems, including PCOS. It affects the pituitary gland. Side effects include nausea, headaches, ovulation pains, and an increased risk of miscarriages and multiple births.

Human menopausal gonadotropin (hMG) is given to women who don't ovulate due to pituitary gland problems. hMG is injected and used directly by ovaries, bypassing the pituitary gland.

Metformin is given to women with insulin-resistance or PCOS. It lowers the levels of male hormones, which are usually high with these problems. Metformin, given orally, helps women ovulate and is often combined with clomidane citrate or FSH.

Bromocriptin is used when women have problems ovulating because of high levels of prolactin, a hormone released when breast-feeding. Side effects include dizziness, nausea, and headaches.

Doxycyline is an oral antibiotic for men and women. Women take it to help their cervical mucus become less hostile. Men take it to decrease the amount of bacteria in their semen. The bacteria attract white blood cells which then attack the sperm.

Other antibiotics can be prescribed for men with infections in the reproductive tract. This doesn't always improve the man's fecundity. There are hormone drugs that men with hormonal imbalances can take [7, 36].

NPT. Natural procreative technology, (NaProTechnology, NPT) is a systematic medical approach for optimizing conceptions in vivo, through the least invasive methods possible. Through a general or family physician, the events that occur during the menstrual cycle are monitored by the couple. Abnormalities commonly found are low amount of cervical mucus, intermenstrual bleeding, short luteal phases, and low levels of estrogen and progesterone. Common interventions for these problems include induction or stimulation of ovulation, medicines to produce more cervical mucus (including vitamin B6, guaifenesin, or antibiodics), and hormone supplements [55].

Psychological. The effectiveness of psychological treatments is undetermined. In a meta-analysis of studies on psychology and infertility, it was concluded that psychological therapy is significant for couples not using ART [21]. There are a variety of psychological interventions. Most effective are cognital behavioral and didatic skill building. One study had 50% birth rate compared to 20% control group [18].

Surgery. Surgical procedures are common interventions for tubal obstructions and excessive endometrial tissue. If there is a small obstruction preventing the egg from entering the fallopian tube, the surgery is typically small and basic. However, if the fallopian tube is blocked on both ends, the surgery is more invasive. Women with either too much endometrial tissue or endometriosis, can have the tissue removed so that the zygote can implant. This is not too invasive and usually consists of cutting the tissue with a laser. Surgery is an option for men who have blockages and varicocele. In the first case, surgery can remove the obstructions, while in the second, the veins are drained [36].

IUI. Intrauterine insemination (IUI), also known as artificial insemination, is when a woman is injected with prepared sperm. Since only the sperm are handled, and the egg is fertilized in vivo, it is not considered an ART procedure. IUI is used for couples with mild male infertility, cervical mucus problems, and unexplained infertility [25].

ART. ART, which stands for assisted reproductive technologies, is a category of fertility treatments where both the sperm and eggs are handled. Common complications are preterm births, multiple births, and low-birth weight. In vitro fertilization, IVF, is the most common ART procedure. As much as 99% of ART procedures are IVF [22]. Mature eggs are taken from the woman's ovaries and placed in a dish with the sperm. The fertilized eggs are then put in the uterus. It is effective when the man has a low sperm count or the path from the ovaries to the uterus is blocked. ZIFT, zygote intrafallopian transfer or tubal embryo transfer, is similar to IVF except that the fertilized egg is placed in the fallopian tube instead of the uterus. This procedure is used when a blockage in the fallopian tube prevents the sperm from fertilizing the egg. ICIS, intracytoplasmic sperm injection, is used when there are serious sperm issues ie., a very low sperm count, failed IVF procedures, or older couples. A single sperm is injected into a mature egg, then placed in the uterus or fallopian tube. When there is a problem with implantation, assisted zona hatching (AZH), in which a small opening is made in the egg to help it hatch, and thus help it implant, is sometimes performed. If IVF

Age	Rate
35 and under	42
35- 37	32
38-40	22
41-42	12
43-44	5

Table 2.4: ART Success Rates by Age [41]

has failed or if the woman has a poor embryo quality, autologous endometrial coculture is an option. In this procedure, the fertilized egg is placed on uterine cells to help it grow. Then it is placed in the uterus, with the hope that implantation will occur more easily. If there is a severe problem with the sperm or eggs, donor eggs or sperm can be used. Occasionally, embryos are frozen using cytopreservation to be used later [22].

Chapter 3. Background Math

3.1 Markov Chains

The model is based on the theory of Markov Chains, which we now review. To create a model, we have to simulate seemingly random events. We call these events random variables. A random variable denoted X, is a variable that is subject to variations due to chance. It is defined on a set Ω of possible outcomes and has a probability distribution which assigns likelihood of happening to each outcome.

Definition 3.1. Let (Ω, A) be a measurable space and $X : \Omega' \to \Omega$ be measurable. X is called a random variable with values in (Ω, A). If $(\Omega, A) = (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, then X is called a real random variable.

Example 3.2. Gambling

Let $\Omega = \{\text{heads,tails}\}$. We define a game in which a gambler tosses a coin. If the coin lands with the heads up, the gambler pays \$1. If a tails is on top when the coins is flipped, the gambler gets \$1. We define the random variable X to be the amount the house pays or receives from a coin flip.

$$X(w) = \begin{cases} -1 & : w = \text{heads} \\ 1 & : w = \text{tails} \end{cases}$$

When using a fair coin, the probability distribution is

$$\rho(y) = \begin{cases} \frac{1}{2} & : y = -1 \\ \frac{1}{2} & : y = 1 \\ 0 & : \text{else} \end{cases}$$

Since most useful models incorporate many random variables, we create a stochastic, or random, process, which is a collection of random variables. It often represents how a value changes over time.

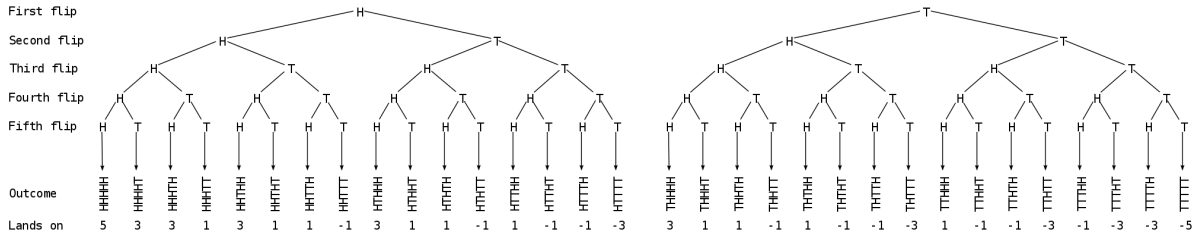

Figure 3.1: The possible outcomes for a random walk of equal probability in five steps

Definition 3.3. Let $I \in \mathbb{R}$ A family of random variables $X = \{X_t | t \in I\}$ on $(\Omega, \mathcal{F}, \mathbf{P})$ with values in (E, ε) is called a stochastic process with index (or time) set I and range E.

Example 3.4. Random Walk

A random walk is a useful example of a stochastic process. It is basically a path that consists of multiple steps. We extend the previous example to a random walk. In this case, the random walk measures how much the gambler has paid or received over a period of time. Since each step consists of a one dollar, the random walk is defined on the integers $\mathbb{Z}$. We assume the gambler starts with \$0.

So if the first toss is heads, then the gambler owes one dollar, moving from zero to owing a dollar. If the next toss is also heads, the gambler then owes two dollars Each toss is a random variable as defined above. The stochastic process is the collection of all of the random variables. Figure 3.1 show the possible outcomes

We can explain the example formally as follows:

Define random variables $\mathbb{Z}_1, \mathbb{Z}_2, \ldots$ that have values either plus one dollar or minus one dollar, with a fifty percent chance of increasing and a fifty percent chance of decreasing. Set $S_0 = 0$ and $S_t = \sum_{j=0}^{t} Z_j$. Then $\{S_t\}$ is a random walk on $\mathbb{Z}$.

One characteristic that stochastic processes can have is called the Markov property. It is named after Andrey Markov, a Russian mathematician from the late eighteen hundreds whose main work was in Markov processes.

The Markov property, a type of memorylessness, means that the future state of the processes only depends on the current state. The example above is memoryless. If the house

must be currently at state \$3, then the only possible immediate future states are \$2 and \$4. Which state the house moves to is independent of how the house reached state \$3.

Definition 3.5. Let E be a polish space with Borel σ-algebra $\mathcal{B}(E), X$ be an E-valued stochastic process, and $\mathcal{F} = \sigma(X)$. X has the Markov property if for all $A \in \mathcal{B}(E)$, and all s, t in I with $s \leq t$, $P[X_t \in A|\mathcal{F}_s] = P[X_t \in A|X_s]$.

This is often expressed as

$$P(X_{n+1} = x|X_1 = x_1, X_2 = x_2, \ldots, X_n = x_n) = P(X_{n+1} = x|X_n = x_n),$$

which means that the probability of moving to state x at time $n + 1$ given that at time 1, the process was in state X_1, and at time 2, it was in state X_2, all the way up until time n when the process was in state X_n, is the same as the probability of moving to state x from state X_n

Definition 3.6. A Markov chain is a random process with the Markov property. It is represented by a stochastic, or transition, matrix.

A Markov chain can also be characterized by the following three properties.

- The outcome of each experiment is one of a set of discrete states.
- The outcome depends only on the present state.
- The transition probabilities remain constant from one transition to the next.

A transition matrix is an n by n matrix where n is the number of states, or values in Ω. Each row represents the transition probabilities from one state to all of the states, represented by columns. Thus the sum of the entries of each row is 1. The probability of moving from state i to state j in n steps is $p_{ij}^n = P(X_n = j|X_0 = 0)$. To find this, we raise the transition matrix to the nth power

Example 3.7. Random Walk Part II

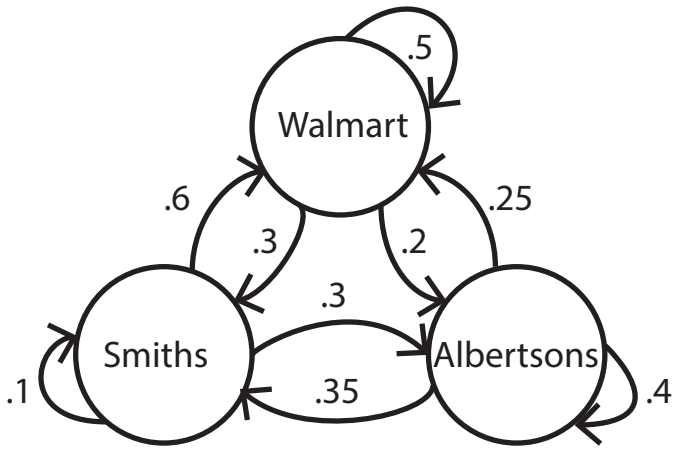

Figure 3.2: Markov Chain for Example 3.8

Expanding the random walk example, we assume that once the gambler owes \$4, he is bankrupt and always owes \$4. We also assume that when then the gambler has \$4, he has attained the maximum amount and always has \$4. Thus once in the gambler is in state 4 or -4, he cannot move. This is called an absorbing state when $p_{ii} = 1$. The transition matrix for this random walk is

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ .5 & 0 & .5 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & .5 & 0 & .5 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & .5 & 0 & .5 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & .5 & 0 & .5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & .5 & 0 & .5 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & .5 & 0 & .5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & .5 & 0 & .5 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

where the rows in descending order and the columns from left to right represent states $4, 3, \ldots, 0, \ldots, -4$.

Example 3.8. Grocery Stores

In a small town, there are three different grocery stores: Walmart, Smiths, and Albertsons. Each week, 50% of people who shop at Walmart will shop at Walmart the following week. 30% will switch to shopping at Smiths and the other 20% will shop at Albertsons. Of the shoppers who went to Smiths, 10% will continue to shop there the following week, 60% will take their business to Walmart, and 30% will shop at Albertsons. Of the shoppers who went to Albertsons, 40% will stick with them, 25% will move to Walmart, and 35% will shop at Smiths the following week. This process is shown in Figure 3.2.

We can model this stochastic process via the transition matrix:

$$T = \begin{bmatrix} .5 & .3 & .2 \\ .6 & .1 & .3 \\ .25 & .35 & .4 \end{bmatrix}$$

The rows represent the current week while the columns represent the following week. The first row represents the shoppers at Walmart. The columns, in order, express Walmart, Smiths and Albertsons. Thus the first entry in the first row is the percentage of shoppers at Walmart who shop at Walmart the following week. The second entry of the same row is the percentage of shoppers at Walmart who shop at Smiths the following week.

Using powers, we can estimate what percentage of shoppers will be buying groceries at Walmart in four months.

$$T^4 = \begin{bmatrix} .4566 & .261 & .282 \\ .455 & .263 & .282 \\ .455 & .263 & .284 \end{bmatrix}$$

This means that in one month, 46% of shoppers will shop at Walmart, 26% will shop at Smiths, and 28% will shop at Albertsons.

Markov processes have been studied extensively and been applied in a wide range of disciplines.

In physics, Markov chains are applied in thermodynamics and Brownian motion. Markov chains have been used in pattern recognition and data compression. Hidden Markov models are a type of Markov process that have developed voice recognition. A famous application is google's page rank formula for returning search engine query results. Other fields include games, music, mathematical biology, statistics, genetics, and even baseball.

3.2 Decision Processes

While Markov chains are extremely useful, they cannot account for the decisions that we make everyday. In the random walk example, the Markov chain assumes that the gambler keeps tossing the coin for a certain amount of time. But the gambler can always choose to stop at any time period. Fortunately, we can extend a Markov chain to a decision process. There are several classes of decision processes, but we will only discuss deterministic Markov decision processes. A Markov Decision Process is a collection of decision epochs, states, transition probabilities, actions, and rewards. The states and transition probabilities define a Markov chain, so a Markov decision process can also be thought of as a Markov chain combined with decisions and rewards.

A decision epoch is the point in time at which a decision is made. We are assuming a discrete time problem, in which time is divided into periods and the number of decision epoch is finite. Define T to be the set of decision epochs. When $|T| < \infty$, we say the decision process has a finite-horizon.

At each time epoch, the system occupies a state $s \in S$, the set of all possible states. For a given s, there is a set of allowable actions, A_s, with $A = \cup_{s \in S} A_s$. We assume that S and A_s do not vary with t since we will not actually know s. We assume that actions are chosen deterministically.

For each a in A_s in state s, there is an associated reward, $r(s, a)$. Sometimes the reward will also depend on the time, denoted $r_t(s, a)$. The system state at the next decision epoch is determined by a transition probability at time t, $p_t(\cdot \mid s, a)$. We assume that $\sum_{j \in s} p_t(j \mid$

$s, a) = 1$.

Definition 3.9. We define a Markov Decision Process to be the collection

$$\{T, S, A_s, p_t(\cdot \mid s, a), r(s, a)\}. \tag{3.1}$$

Example 3.10. Cake Eating

A classic deterministic Markov decision process problem is called the "cake eating" problem. Say that you have a cake and you want to maximize your enjoyment from eating the cake before it goes stale. How much cake should you eat every day? If you eat the whole cake the first day, chances are you'll feel pretty sick and you won't be able to enjoy some the next day. Satisfaction from eating the cake decreases with each bite, so eating a lot of cake on one day won't be as good as eating a little bit each day.

This example falls under the broader category of sequential allocation models, in which there are M units of resources available for use over N time periods. During time period t, x_t units are consumed. $f(x_1, x_2, \ldots, x_N)$ is the utility function. The goal of the model is to maximize utility. When f is separable,

$$f(x_1, \ldots, x_N) = \sum_{t=1}^{N} g_t(x_t) \tag{3.2}$$

the problem is easily formulated as a Markov Decision Process.

Formulation

Decision Epochs: $T = \{1, \ldots, N\}, N < \infty$

States: $S = [0, M], M < \infty$

Actions: $A_s = [0, s], 0 \leq s \leq M$

Rewards: $r_t(s, a) = g_t(a), s \in S, a \in A_s, t = 1, \ldots, N-1$ and $r_N(s) = g_N(s), s \in S$

Transition Probabilities:

$$p_t(j|s, a) = \begin{cases} 1 & : j = s - a \\ 0 & : \text{else} \end{cases} \quad s \in S, a \in A_s, t = 1, \ldots, N-1$$

Note that it is because the transition probabilities are binary that we have a deterministic Markov process.

3.3 Backward Induction

To solve a finite-horizon, discrete-time Markov decision process, we use the Backwards Induction Algorithm, which is an efficient way to find the optimal policy.

The idea is to choose the action that maximizes the expected reward at the last state. Moving forward in time, we calculate the total expected reward from the possible actions, based on the expected reward in the last state, and choose the action that maximizes the total expected reward.

We use the notation from Example 3.10 and introduce two new terms.

$u_t^*(s)$ is the total expected reward at time t for state s.

$A_{s,t}^*$ is the action that should be chosen at time t for state s.

Backwards Induction Algorithm

1: Set $t = N$ and

$$u_N^*(s_N) = r_N(s_N), \forall s_N \in S \tag{3.3}$$

2: Set $t = t - 1$ and calculate

$$u_t^*(s_t) = \max_{a \in A_{s_t}} \{r_t(s_t, a) + \sum_{j \in S} p_t(j \mid s_t, a) u_{t+1}^*(j)\} \tag{3.4}$$

Set

$$A_{s_t,t}^* = \text{argmax}_{a \in A_s} \{r_t(s_t, a) + \sum_{j \in S} p_t(j \mid s_t, a) u_{t+1}^*(j)\} \tag{3.5}$$

3: If t = 1, stop. Else, return to step 2.

Theorem 3.11 ([51], p. 92-93)**.** *The backwards induction algorithm has the following properties:*

Suppose that $u_t^*, t = 1, \ldots, N$ *and* $A_{s_t,t}^*, t = 1, \ldots, N-1$ *satisfy equation 3.4 and equation 3.5 in the backwards induction algorithm. Let* h_t *be the history of states and actions up until time* t*.*

$$h_t = (s_1, a_1, s_2, \ldots, a_{t-1}, s_t) \tag{3.6}$$

π *is a policy comprised of decisions, or actions,* $(d_1, d_2, \ldots d_{N-1})$ *and* π^* *is the optimal policy. Then*

a) for $t = 1, \ldots, N$ *and* $h_t = (h_{t-1}, a_{t-1}, s_t)$

$$u_t^*(s_t) = \sup_{\pi \in \Pi^{HR}} u_t^{\pi}, s_t \in S \tag{3.7}$$

b) Let $d_t^*(s_t) \in A_{s_t,t}^* \forall s_t \in S, t = 1, \ldots, N-1$ *and let* $\pi^* = (d_1^*, \ldots, d_{N-1}^*)$*. Then* $\pi^* \in \Pi^{MD}$ *is optimal and satisfies*

$$u_t^{\pi^*}(s_t) = u_t^*(s_t), s_t \in S, t = 1, \ldots, N \tag{3.8}$$

This theorem means that

a) for $t = 1, \ldots, N-1$, the backwards induction algorithm finds the sets $A_{s_t,t}^*$ which contain all actions in A_{s_t} which attain the
maximum in step 2.

b) the backwards induction algorithm, evaluates any policy that selects an action in $A_{s_t,t}^*$ for each $s_t \in S$ for all time periods until the end.

c) the algorithm computes the expected total reward for the entire decision-making horizon, and from each period to the end of the horizon for any optimal policy.

3.4 Multi-Armed Bandit Problem

One class of Markov decision processes is the multi-armed bandit problem, or N-armed bandit problem. They are equivalent to one-state Markov decision processes. The name

comes from the classic gambling example; a gambler has a row of slot machines, also known as one-armed bandits. Each machine has its own reward distribution. Which machines, in which order, should the gambler play to maximize his total expected reward? Since the gambler doesn't know the machines' distributions, the policy has to balance exploitation, choosing the arm with the highest probability of a reward, with exploration, learning more about a distribution to gain greater reward in the future.

Formulation

Decision Epoch: $T = \{1, 2, \cdot, N\}$

States: $S = S_1 \times S_2 \times \cdots \times S_k$

Actions: $r((s_1, s_2, \cdots s_k), a) = r_i(s_i)$ if $a = i$

Transition Probabilities:

$$p((u_1, \cdots, u_k) \mid (s_1, \cdots, s_k), i) = \begin{cases} p(j_i \mid s_i) & : u_i = j_i, u_m = s_m \text{ for } m \neq i \\ 0 & : u_m \neq s_m \text{ for some } m \neq i \end{cases}$$

Multi-armed bandit problems have been used to simulate clinical trials. Using a simulation can be helpful when a laboratory cannot test enough patients.

The typical clinical trial problem is: there are n arms/treatments. These arms may be pulled, using the gambling example terminology, or given to a patient, repeatedly in any order. Each pull results in a success or a failure. The sequence of successes and failures for a given arm forms a Bernoulli process with unknown success probability θ. A success results in a reward of the form a^{t-1}. $a \in (0, 1)$ is known as the discount factor. A failed pull yields no reward. So assume a patient is given treatment i. If it succeeds, their reward of being healed will be affected by time. One way to look at it is that the patient will spend more money over a long period of time, reducing his total reward. If the treatment fails, obviously the patient isn't healed and has no reward. At the initial time period, the probability for each arm has a beta prior distribution. We assume independence of arms. This allows us to us Bayes' Theorem as the arms are pulled to change the prior distributions to posterior

distributions. Because the class of beta distributions is closed under Bernoulli sampling, the posterior distributions are also beta distributions, which means that we can continue updating the prior and posterior distributions for multiple pulls.

One solution is pull the arm for which the expected value of θ is largest. Arms $i = 1, 2, \ldots, n$ with probabilities of success θ_i and distributions $Beta(\alpha_i, \beta_i)$ have means $\frac{\alpha_i}{\alpha_i+\beta_i}$. Thus we'd pull arm j with the property that

$$\frac{\alpha_j}{\alpha_j + \beta_j} = \max \frac{\alpha_i}{\alpha_i + \beta_i} \tag{3.9}$$

While in general, this policy is quite good, it is not optimal. It does not address the problem of which arm to choose when the mean of two arms are equal and achieve the maximum mean.

$$\frac{\alpha_j}{\alpha_j + \beta_j} = \frac{\alpha_k}{\alpha_k + \beta_k} = \max_i \frac{\alpha_i}{\alpha_i + \beta_i} \tag{3.10}$$

This lack of foresight, known both mathematically and literally, is referred to as myopic, and can be solved through a recursive dynamic programming equation.

$$\begin{aligned} R(\alpha_1, \beta_1, \alpha_2, \beta_2) = \max\{ & \frac{\alpha_1}{\alpha_1 + \beta_1}[1 + aR(\alpha_1 + 1, \beta_1, \alpha_2, \beta_2)] && (3.11) \\ & + \frac{\beta_1}{\alpha_1 + \beta_1} aR(\alpha_1, \beta_1 + 1, \alpha_2, \beta_2), && (3.12) \\ & \frac{\alpha_2}{\alpha_2 + \beta_2}[1 + aR(\alpha_1, \beta_1, \alpha_2 + 1, \beta_2)], && (3.13) \\ & + \frac{\beta_2}{\alpha_2 + \beta_2} aR(\alpha_1, \beta_1, \alpha_2, \beta_2 + 1)\} && (3.14) \\ & && (3.15) \end{aligned}$$

It can be shown that the above equation converges to an optimal reward function [27].

However, it is computationally expensive, requiring

$$\frac{(N-1)!}{(2n)!(N-2n-1)!} \tag{3.16}$$

calculations and

$$\frac{(N-1)!}{(2n-1)!(N-2n)!} \tag{3.17}$$

of storage. This becomes very difficult computationally for $n > 3$, so we apply the index theorem.

Definition 3.12. Index policy: There exists a real-valued index $v(B_i, \psi_i)$ which may be computed for each bandit process separately as a function of its current state only. The optimal policy is to choose the arm with the greatest index.

Theorem 3.13 ([27], p. 27)**.** *Gittins Index Theorem: A policy for a simple family of alternative bandit processes is optimal if it is an index policy with respect to* $v(B_1, \cdot), ...v(B_n, \cdot)$.

Using the Gittins Index Theorem, we assume that $n = 2$, and arm 1 has success probability θ which has a prior distribution $Beta(\alpha, \beta)$. We also assume that arm 2 has a known success probability p. The dynamic programming recurrence equation is

$$R(\alpha, \beta, p) = \max\left\{\frac{p}{1-a}, \frac{\alpha}{\alpha+\beta}[1 + aR(\alpha+1, \beta, p)] + \frac{\beta}{\alpha+\beta}aR(\alpha, \beta+1, p)\right\} \tag{3.18}$$

This equation is equivalent to the previous when $\frac{p}{1-a}$ is replaced with $p + aR(\alpha, \beta, p)$ for a given p. $\frac{p}{1-a}$ is the total expected value from a policy which always pulls arm 2, since at time t, the expected reward is $a^{t-1}p$, which is a geometric series. If it is optimal to pull arm 2 at any time, then it must be optimal to pull arm 2 every time after that, since there is no new information about arm 1, keeping the second part of equation 3.18 the same.

We can solve equation 3.18 iteratively starting with $\alpha + \beta = N$. When

$$\frac{p}{1-a} = \frac{\alpha}{\alpha+\beta}[1 + aR(\alpha+1, \beta, p)] + \frac{\beta}{\alpha+\beta}aR(\alpha, \beta+1, p) \tag{3.19}$$

the corresponding α and β describe the arms with beta prior distributions whose index value is equal to p. Index values can be then calculated for all beta priors by solving equation 3.18 for a sufficiently fine grid or p-values. Then the optimal policy for any n is to pull the arm with the highest index value. Only $\frac{1}{2}(N-1)(N-2)$ calculations are required to create the Gittins index and the storage requirement is $N-1$, which makes the index much easier to compute than equation 3.1.

Example 3.14. Bandit Index for N = 4

$$\begin{aligned} \text{Let } N &= 4 \\ \text{p-values} &= \{.01, .02, \ldots, .98, .99\} \\ \theta &= .4 \\ a &= .09 \\ p &= .35 \\ \frac{p}{1-a} &= \frac{.35}{1-.09} = 3.5 \end{aligned}$$

To initialize, or compute $R(\alpha, \beta, p)$ for $\alpha + \beta = N$, we set the second part of equation 3.18 equal to $\frac{p'}{1-a}$ where $p' = \frac{\alpha}{\alpha+\beta}$ since $\frac{\alpha}{\alpha+\beta}$ is the mean of the distribution so we have the same geometric series trait as we did with p. So

$$R(0,4,.35) = \max\{3.5, 0\} = 3.5$$

$$R(1,3,.35) = \max\{3.5, 2.5\} = 3.5$$

$$R(2,2,.35) = \max\{3.5, 5\} = 5$$

$$R(3,1,.35) = \max\{3.5, 7.5\} = 7.5$$

$$R(4,0,.35) = \max\{3.5, 10\} = 10$$

Then we calculate the recursive equation for $\alpha + \beta = 3$

$$R(0,3,.35) = \max\{3.5, \frac{0}{3}[1 + .9(3.5)] + \frac{3}{3}.9(3.5)\} = \max\{3.5, 3.15\} = 3.5$$
$$R(1,2,.35) = \max\{3.5, \frac{1}{3}[1 + .9(5)] + \frac{2}{3}.9(3.5\} = \max\{3.5, 3.33\} = 3.93$$
$$R(2,1,.35) = \max\{3.5, \frac{2}{3}[1 + .9(7.5)] + \frac{1}{3}.9(5)\} = \max\{3.56.66\} = 6.66$$
$$R(3,0,.35) = \max\{3.5, \frac{3}{3}[1 + .9(10)] + \frac{0}{3}.97.5\} = \max(3.5, 10\} = 10$$

And for $\alpha + \beta = 2$.

$$R(0,2,.35) = \max\{.35, \frac{0}{2}[1 + .9(3.93)] + \frac{2}{2}.9(3.5)\} = \max\{3.5, 3.15\} = 3.5$$
$$R(1,1,.35) = \max\{.35, \frac{1}{2}[1 + .9(6.66)] + \frac{1}{2}.9(3.93)\} = \max\{3.5, 5.07\} = 5.27$$
$$R(2,0,.35) = \max\{.35, \frac{2}{2}[1 + .9(10)] + \frac{0}{2}.9(6.66)\} = \max\{3.5, 10\} = 10$$

We perform the same calculations for all p values. For each combination of α and β, we determine the p-value for which $R(\alpha, \beta, p)$ has equality in the two entries. So for each

Reward for α and β	p-value
R(0,4)	0
R(1,3)	.25
R(2,2)	.5
R(3,1)	.75
R(4,0)	1
R(0,3)	.32
R(1,2)	.46
R(2,1)	.74
R(3,0)	1
R(0,2)	.32
R(1,1)	.68
R(2,0)	1

Table 3.1: Bandit Index for $N = 4$, $a = .9$

combination of α, β, the index value is the corresponding p-value, $R(\alpha, \beta)$.

Chapter 4. Model

The model consists of four sections; creating the infertility population, creating the bandit index, creating the backwards induction policy, and putting it all together. These sections outline the method. The code has not been included, but it is available on request.

4.1 Modeling the Infertility Population

The fertility process can be described by a randomized Markov chain. The states are the different stages of conception.

To generate the infertile population, we created couples with specific health categories and then simulated a Markov chain to see if they would be classified as "infertile".

4.1.1 Creating Couples. Each couple consisted of the woman's age, health categories, and their transition probabilities.

Since age is an important factor in the woman's ability to get pregnant, each couple was randomly assigned an age between 18 and 48. While most 15 and 48-year-old women are not trying to get pregnant, these ages were included for completeness and the rare instances in which the woman is attempting to conceive. Male ages were not taken into account, even though that can have an effect on infertility.

Each couple was assigned a health category. The three main categories are healthy, unexplained infertility, and explained infertility. A couple can only be one of these categories. The explained infertility is broken up into two further categories, male and female. A couple can have both male and female infertiltiy, so it is possible to be in both subcategories. Male infertility is broken up into azoospermia, oligospermia, and other infertility for males. A couple can be in only one male category. The female infertility is split up into anovulation, endometriosis, cervical problems, tubal and uterine problems, and other female infertility. A woman can have multiple problems, so the possible female infertility category combinations used are

- anovulation
- anovulation and endometriosis
- anovulation and cervical
- anovulation and tubal
- endometriosis
- endometriosis and cervical
- endometriosis and tubal
- cervical
- cervical and tubal
- tubal
- other

Thus there are 35 total health categories. These categories were chosen based on the diagnosis used in the studies "The Spontaneous pregnancy prognosis in untreated subfertile couples," "The prognosis for live birth among untreated couples," and "Population study of causes, treatment, and outcome of infertility " [29, 53, 16]. These studies surveyed the infertile populations, calculated the percentages of couples with different diagnosis, and calculated the live birth rates, without interventions, for these couples. Table 4.1.1 shows which categories were included in which papers. We included mucus failure/dysfunction in the tubal category, since the mucus is found in the Fallopian tubes. We chose to include other male and other female infertility to create a more complete model. Some of the possible causes discussed in Chapter 2 are not accounted for in the categories, including sperm dysfunction, lifestyle choices, antibodies, and infections.

Diagnosis	Spontaneous [53]	Prognosis [16]	Population Study [29]
unexplained	yes	yes	yes
ovulation defect	yes	yes	yes
oligospermia	yes	yes	yes
azoospermia	yes	yes	yes
tubal defect	yes	yes	yes
endometriosis	yes	yes	yes
cervical problem	yes	no	no
mucus failure/dysfunction	no	yes	no

Table 4.1: Diagnosis groups found in studies

Here are some examples of health categories generated by the code.

$[1, 0, [0, 0, 0], [0, 0, 0, 0, 0]]$ healthy

$[0, 1, [0, 0, 0], [0, 0, 0, 0, 0]]$ unexplained

$[0, 0, [0, 1, 0], [1, 0, 0, 0, 0]]$ oligospermia and anovulation

$[0, 0, [0, 0, 0], [0, 1, 0, 1, 0]]$ endometriosis and tubal

Based on the age of the woman and the health category of the couple, transition probabilities were assigned the couple.

We used ten states in the random Markov chain, see Figure 4.1: not pregnant(between cycles), ovulation, fertilization, implantation, 1st trimester, miscarriage, 2nd trimester, 3rd trimester, stillbith, and livebirth. Transition probabilities for each health category were assigned and then adjusted for age. We assume that live birth is an absorbing state since the goal of the simulation is to achieve a live birth. This may not always be true in practice since many couples desire more than one child. Rows 6 and 9 represent recovery time for miscarriages and stillbirths. Woman who experience these events are typically not fertile again for a short period of time as their body readjusts to their normal cycle.

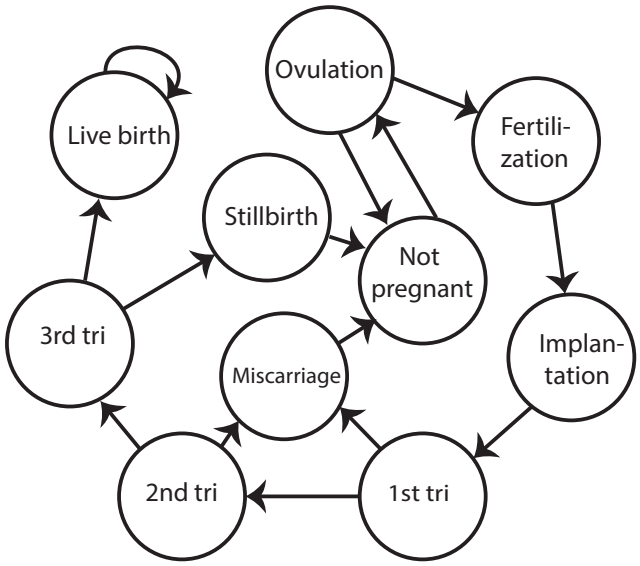

Figure 4.1: Markov Chain for Conception in Model

The transition matrix has the form

$$\begin{bmatrix}
\text{1-ovulate} & \text{ovulate} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\text{1-fertilize} & 0 & \text{fertilize} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\text{1-implant} & 0 & 0 & \text{implant} & 0 & 0 & 0 & 0 & 0 & 0 \\
\text{1-1st} & 0 & 0 & 0 & \text{1st} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & \text{miscarriage} & \text{2nd} & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0. \\
0 & 0 & 0 & 0 & 0 & \text{miscarriage} & 0 & \text{3rd} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \text{stillbirth} & \text{live birth} \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}$$

The probabilities in the transition matrix were chosen such that taking the power agreed

with data from the studies by Collins and Snick [16, 53].

So in n time periods, a couple should have the same chance of having a live birth as the studies showed for the equivalent number of months. To calculate this, we found the mean number of months couples in each category were included in the study. We adjusted this time to correspond to the time periods for the transition matrix, since the time to pregnancy in the model is four periods, while in life, it is only one. Similarly, the trimesters are three months but only one time period each. Assuming that each couple is considered infertile and has been attempting to become pregnant for a year, to calculate the total number of time periods, We added sixteen time periods. The final numbers used are the means of the two studies. Table 4.2 shows the live birth rate and the average number of months couples were included in the study, as well as the chance of having a live birth in the given number of time periods from the model.

For example, the live birth rate for couples with unexplained infertility is 38.6%. The average number of months each couple was included in the study was 15.8. 16 months is equivalent to 14 time periods, so the number of time periods is 30. Then the transition matrix T for unexplained infertility satisfies $T^{30}[0][9] = .38$, or the probability of moving to a live birth state from the not pregnant state in 30 time periods is .38. While this overall probability is satisfied, there are other combinations of the individual probabilities in T that could also satisfy this constraint.

The ovulation defects, other male infertility, and other female infertility categories had to be adjusted. The average number of months women with ovulation defects were included in the studies was only 7.1 months. This number doesn't make sense when calculating live birth rates, so it is assumed that some couples dropped out fairly early on. We used 28 months as the number of time periods, since it was the mode of the other time periods. Since other male and other female infertility were not included, We chose to base them off of the unexplained rates. This may have made them somewhat redundant, but the fact that they could be paired with other infertility causes and that the individual probabilities were

different justified their inclusion in the model.

Diagnosis	Studies live birth rate	Studies average # of months	Model chance of live birth	Model # of time periods
unexplained	38.6	15.8	.38	30
oligospermia	19.4	13.6	.20	28
azoospermia	0	18	0	33
tubal defect	5.9	13.4	.06	28
endometriosis	9.42	13.1	.09	28
cervical problem	11.94	18.6	.18	33
ovulation defect	8.8	7.1	.09	28
female other	NA	NA	.38	30
male other	NA	NA	.36	30

Table 4.2: Transition Statistics

This is the transition matrix for a thirty-two-year-old woman and a man with oligospermia.

$$\begin{bmatrix}
0.6856 & 0.3144 & 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\
0.968 & 0. & 0.032 & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\
0.4 & 0. & 0. & .6 & 0. & 0. & 0. & 0. & 0. & 0. \\
0.5 & 0. & 0. & 0. & 0.5 & 0. & 0. & 0. & 0. & 0. \\
0. & 0. & 0. & 0. & 0. & 0.2 & 0.8 & 0. & 0. & 0. \\
1. & 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\
0. & 0. & 0. & 0. & 0. & 0.04 & 0. & 0.96 & 0. & 0. \\
0. & 0. & 0. & 0. & 0. & 0. & 0. & 0. & 01. & 0.99 \\
1. & 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0. & 0. \\
0. & 0. & 0. & 0. & 0. & 0. & 0.1 & 0. & 0. & 1.
\end{bmatrix}$$

So this woman has a 31.44% chance of ovulating normally, moving from state 1 to state 2. Once in state 2, there is a 3.2% chance of the egg being fertilized, moving the woman into

Diagnosis	% in studies	% in model
Unexplained	27.86	23.4
Ovulation Defect	24.53	22.8
Endometriosis	7.4	7.2
Cervical Defect	15.3	15.6
Tubal Defect	17.76	17.1
Azoospermia	5.85	5.6
Oligospermia	20.16	21.3

Table 4.3: Percentage of infertile couples with different diagnosis

state 3 and a 96.8% chance of the egg dissolving, putting the woman back into state 1.

To determine if a couple was infertile, the couple started in the first state. A random number r was generated. If r was less than (1 - ovulation probability), the woman stayed in state 1. Otherwise, she moved into state 2. This was done for each time period. If a couple failed to be in state 7 or 9, the third trimester or live birth, the couple was classified as infertile. We note that the percentages of infertile couples in each category, found in Table 4.3, are close to the averages found in [53, 29, 16] when tested for 10000 infertile couples.

4.2 Backwards Induction

The Markov decision process for the amount of resources consumed is a resource allocation problem, similar to that of Example 3.10. The formulation for this problem is

- P: length of time period, estimated to be 4 weeks
- N: the number of time epochs, with a max of 36
- T: the set of time epochs
- t: the current time epoch
- S: the set of all states.
- $s \in S$: each state is a triple, consisting of the time remaining, the budget remaining, and a constraint tuple. The maximum time remaining we used is 36 months, or three

years. The maximum budget we used is $60000. Both of these variables can be easily changed, but increasing them increases the amount of time and storage required to create the index. The constraint array is a list of maximum number of constrained actions available. The model has the following interventions constrained, ART, IUI, male surgery, and female surgery. There were two reasons for implementing a constraint array. First, it reduces the amount of time needed to calculate the optimal policy. Second, most men and women won't have multiple surgeries. If a surgery doesn't work the first time, chances are, it won't work the second time. Similarly with ART and IUI, after a certain number of tries, the percentage of success is so small, that it is not worth the couple trying again. We allowed a maximum of 3 attempts for ART, 6 attempts for IUI and 1 attempt for male and female surgeries. However, these number can be easily changed as they are an input into the model. Constraining other interventions will involve changing the backwards induction module.

- A: the set of allowable actions. We assume that $A_s = A$ so each action is allowed at any state. The actions are the possible interventions. The interventions chosen for this project are ART, IUI, male surgery, male hormones, female surgery, estrogen, progesterone, and no treatment. These were chosen because they are the most common intervention categories. It is possible to split them up further, there are several types of ART, but due to feasibility constraints, it was decided that 8 actions was sufficient.

- $r(s, a)$: The reward for choosing an intervention based on the state. The reward consists of the change in state and and the probability of having a live birth after the intervention. The numbers for this category were the most difficult to choose because finding information on them was hard.

intervention	cost	time needed	probability of success
ART	12400	1	.30
IUI	1200	1	.13
male surgery	5000	6	.2
male hormones	200	3	.1
female surgery	4000	5	.1
estrogen	200	3	.1
progesterone	200	3	.9
none	0	1	.05

Table 4.4: The rewards for each treatment

ART: The percentages of success vary widely with ART, ranging from 5% to 45% [23, 25, 10]. Costs range depending on the services used, but are usually over $10000, with an average of $12400 per cycle [12, 52, 47]. Sometimes at least one month is recommended between cycles, but it is not required, so we chose to use one period [23].

IUI: The success rate is between 5 and 20% so we used the average [58, 11]. Costs range from $300 to $4000 including ultrasounds, hormones, triggers and blood work. We estimated $1500 per cycle [58, 30]. There is no recovery time for IUI so the time cost is one period. To speed up the induction program, we used a budget interval of $200. Thus every budget reward is a multiple of that.

For example, if a woman's probability of having a live birth is .03 and she does IVF, her reward will be .3, which is the probability of having a live birth with IVF. It takes one time period to try IVF, $12500, and $[-1, 0, 0, 0]$ is the change in the constraint array. Thus if the current state is $[6, 20000, [2, 2, 1, 0]]$, the reward is $[4, 7500, [2, 3, 1, 0]]$.

- $p(j|s, a)$: the probability of being in state j after choosing action a in state s is 1 if

$r(s,a) = j$ and 0 else.

$$p(j|s,a) = \begin{cases} 1 & : r(s,a) = j \\ 0 & : else \end{cases}$$

4.3 Bandit Model

The bandit model uses the Gittins index in Section 3.4 The discount factor a, number of time periods N, number of p-values, and a threshold value used to speed up the algorithm are all inputs to creating the index, so they can easily be changed. We used $a = .9$. $N = 1000$, number of p-values = 100, and threshold .05. We also started with $\alpha = 1$ and $\beta = 1$, so we ignored the cases with $\alpha = 0$ and $\beta = 0$.

4.4 Putting it all together

As customary with bandit based clinical trials, each intervention was given an initial starting success and failure rate of 1. However, because in this simulation, most couples would try multiple interventions, each couple was also assigned a individual starting success and failure rate of 1. While the population success and failure rates carried over to each couple, the individual values did not. This accounted for a couple's propensity for certain treatments while still accounting for the general success of the treatment. For example, if a woman isn't producing eggs, IVF isn't going to help her become pregnant, even though IVF will have a higher population success rate. For each couple, we randomly assigned a time and monetary budget. For each time period, the optimal policy generated by their budget was calculated. So we have three optimal actions based on different criteria. To decide which action is optimal, we turned the policy into an index by adding the locations of an action. So for policy

$$[0, 1, 1, 3, 6, 6, 7, 7, 7]$$

the index values of arms $0, \ldots, 7$ are

$$[0, 3, 0, 3, 0, 0, 9, 21].$$

We also divided each arm index by the sum of the indices of the policy plus one so that each index value was a percent. This gave shorter policies more weight, since near the end of the allotted time periods, we want to exploit instead of exploring. For policy $[0, 1, 1, 3, 6, 6, 7, 7, 7]$, I'd divide by 45 to get

$$[0.0, 0.02, 0.0, 0.02, 0.0, 0.0, 0.2, 0.47].$$

We then weighted each index to create a final index value. The population bandit index was multiplied by 0.3. The individual bandit index was multiplied by 0.4, and the policy index was multiplied by 0.2. These weights were chosen so that each index value was less than one, the individual bandit index was worth more than the population bandit index (since we are testing individuals, not populations), and that the policy index was the smallest; the policy index might be best financially, but it doesn't account for how good the treatments are. Then the arm with the highest index was checked for feasibility. If it was allowable, that arm was chosen. If not, then the arm with the next highest index was checked for feasibility, and so on, until an allowable arm was reached.

Example 4.1. Choosing the Optimal Arm

Assume

$$\begin{aligned}
\text{population bandit value} &= [0.7, 0.21, 0.38, 0.38, 0.38, 0.32, 0.32, 0.15] \\
\text{individual bandit value} &= [0.7, 0.38, 0.7, 0.7, 0.7, 0.5, 0.5, 0.5] \\
\text{policy} &= [1, 1, 1] \\
\text{policyindex} &= [0.0, 0.86, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0]
\end{aligned}$$

The weighted index values are

$$.3 * [0.7, 0.21, 0.38, 0.38, 0.38, 0.32, 0.32, 0.15] \tag{4.1}$$

$$+.4 * [0.7, 0.38, 0.7, 0.7, 0.7, 0.5, 0.5, 0.5] \tag{4.2}$$

$$+.2 * [0.0, 0.86, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0] \tag{4.3}$$

$$=[.49, .387, .396, .396, .396, .296, .296, .245] \tag{4.4}$$

The order of the indices from highest to lowest are

$$[0, 4, 3, 2, 1, 6, 5, 7]$$

so action 0 is optimal.

Chapter 5. Results

Determining the effectiveness of the model was difficult, since there are no benchmarks or models to compare it to. Intuitively, the results make sense because there is a balance of all the interventions, implying that the model does explore. More costly interventions occur near the end of the time frame, indicating that the model is exploiting resources at the end, when there is not time to explore.

We believe this model to be valid because it matches the high end of success rates in studies for IVF and IUI and sterility estimates. As these are the only statistics we have, these are only numerical comparisons we can make.

age group	% that got pregnant	% of population who are sterile
under 30	.73	.04
under 35	.67	.05
under 40	.6	.06
under 48	.13	.13

Table 5.1: Percent of population that was sterile, assuming 15% infertility rate

In a test case of 1000 infertile couples, IUI was successful 17% of the time. This seems a good estimate, since the rate of success is between 5 and 20% and the population was skewed towards younger women. The ART success rates were similar to the UK 2010 rates and below the U.S. rates.

age group	% from UK	% from US	% from model
under 35	32.2	42	34
under 40	24.25	32	32
under 48	6.83	8.6	7.9

Table 5.2: Percent of women who achieved birth from ART, [10, 41]

As was mentioned earlier, this model is designed to balance exploration with exploitation.

The table below shows how many times each intervention was used during a simulation. Each intervention was used by at least 14% of the couples. Interventions with higher success rates were attempted more often, which we would expect. More costly interventions occur more frequently as time runs out, showing that exploration occurs near the beginning of the process and exploitation occurs near the end.

intervention	number of times tried
ART	558
IUI	499
Male Surgery	150
Male Hormones	235
Female Surgery	86
Estrogen	382
Progesterone	142
None	2172

Table 5.3: Number of times each intervention is attempted for 610 infertile couples

Most of the couples who became pregnant in the test cases did so within a year. The average time to pregnancy for non-sterile couples was 11.2 months. Since couples can try for years before becoming pregnant, having a shorter diagnosis period indicates a valid method of diagnosing.

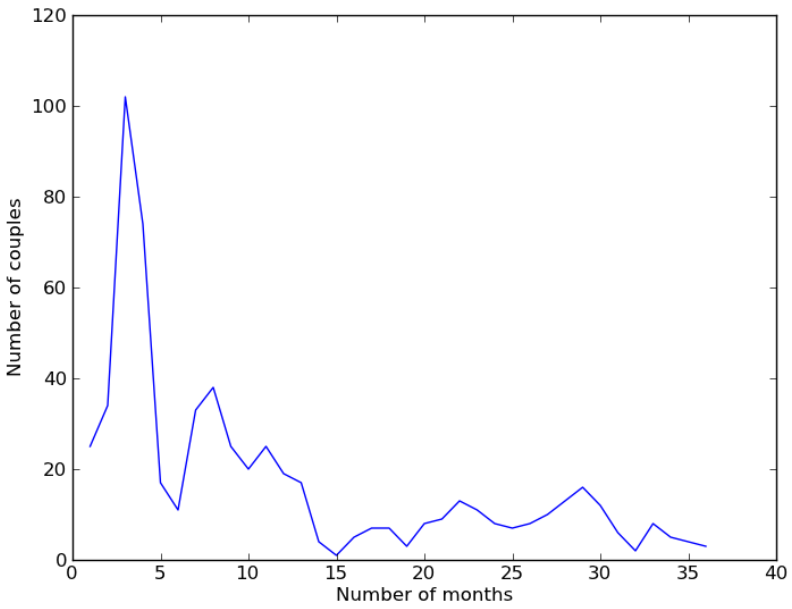

Figure 5.1: Number of couples that became pregnant after a number of months

In 82% of the successful pregnancies, the treatment that worked was appropriate for the cause of infertility. While this number is good, we would like it to be higher. However, as the model is tuned, we believe this number will increase, proving to be a helpful tool in diagnosing and treating infertility.

Chapter 6. Conclusion

This model presents a method for diagnosing infertility via Markov decision processes and multi-armed bandit indices. Due to the complexity of the problem, there are still many ways in which we can improve the model. However, for a proof of concept, the results look promising.

6.1 Future work

There are several ways this model could be improved.

We think that the most important extension is to account for other factors in the reward. First, we could include psychological factors. Since infertility is such a sensitive and emotional subject, this is quite important but is not taken into account in the model. Trying more expensive and invasive procedures typically causes more stress and anxiety, especially when they fail, which can affect how often couples want to attempt them. The side effects of rewards could be included. For instance, with ART, there is a high probability of multiple births. This can be a factor in whether or not couples should choose this procedure. Second, an important aspect of rewards that is not included in the model is how age affects the rewards. Since the probability of a successful treatment is heavily determined by age, this is crucial in determining which intervention a couple should use. Another important aspect of rewards is that interventions can change the probabilities of other interventions working. For instance, female surgery can leave scar tissue, making it harder to implant. Accounting for these downstream probabilities can be really important. Going along with this, the classical multi-armed bandit problem used in this model requires independence of arms. However, as noted above, interventions can affect other interventions. It would be interesting to consider what happens when the arms are not independent.

While the woman's age is taken account, the man's age can have an effect on infertility. This is not nearly as drastic or important as the woman's age, but it is something that should

be considered. Another assumption of the model is that the couple only cares about giving birth once. This may or may not be true for the couple, so accounting for the number of children the couple would like to have would be a nice extension of the model.

Another interesting extension is to factor in medical tests. When a couple first sees a doctor for infertility, there are routine tests that can be done, such as semen count and hormone levels. If a woman has low estrogen, Then the decision process can be modified to only consider treatments applicable to the test results. For instance, if a woman has PCOS, the couple wouldn't need to consider male hormones or surgery.

Finally, the model can always be extended by adding more actions, either in new categories, or creating more specific arms.

Bibliography

[1] *The PDR Family Guide to Women's Health and Prescription Drugs*. Medical Economics, 1994.

[2] Mark Hamilton Abha Maheshwari1 and Siladitya Bhattacharya. Effect of female age on the diagnostic categories of infertility. 2008.

[3] American Pregnancy Association. Endometriosis. `http://www.americanpregnancy.org/womenshealth/endometriosis.html`, Nov 2011.

[4] American Pregnancy Association. Miscarriage. `http://www.americanpregnancy.org/pregnancycomplications/miscarriage.html`, Nov 2011.

[5] American Pregnancy Association. Female fertility testing. `http://www.americanpregnancy.org/infertility/femalefertilitytesting.html`, March 2012.

[6] American Pregnancy Association. Female infertility. `http://www.americanpregnancy.org/infertility/femaleinfertility.html`, Feb 2012.

[7] American Pregnancy Association. Infertility medications. `http://americanpregnancy.org/infertility/infertilitymedications.html`, Feb 2012.

[8] American Pregnancy Association. Male fertility testing. `http://www.americanpregnancy.org/infertility/malefertilitytesting.html`, April 2012.

[9] American Pregnancy Association. Understanding ovulation. `http://www.americanpregnancy.org/gettingpregnant/understandingovulation.html`, 2013.

[10] Human Fertilisation Embryology Authority. Latest uk ivf figures: 2010 and 2011. `http://www.hfea.gov.uk/ivf-figures-2006.html#1278`, Feb 2013.

[11] BabyCenter. Fertility treatment: Artificial insemination (iui). `http://www.babycenter.com/0_fertility-treatment-artificial-insemination-iui_4092.bc`, 2013.

[12] BabyCenter. Fertility treatment: In vitro fertiliation (ivf). `http://www.babycenter.com/0_fertility-treatment-in-vitro-fertilization-ivf_4094.bc?page=2#articlesection4`, 2013.

[13] Jane Bates. *Practical Gynaecological Ultrasound*. Cambridge University Press, 2007.

[14] Modi Deepak et al Chatterjee, Suparna. Screening for foxl2 gene mutations in women with premature ovarian failure: an indian experience. *Reproductive BioMedicine Online*, 15(5):554–560, 2007.

[15] Lisa Church. *Hope is Like the Sun*. HopeXchange Publishing, 2004.

[16] Burrows Elizabeth A. Willan Andrew R. Collins, John A. The prognosis for live birth among untreated infertile couples.

[17] Von Eckardstein S et al Cooper TG, Noonan E. World health organization reference values for human semen characteristics. 2010.

[18] Domar Alice D Cousineau, Tara M. The benefit of psychological interventions for female infertility patients.

[19] Norine Dworkin-McDaniel. The invisible pain of infertility. http://www.redbookmag.com/health-wellness/advice/dealing-with-infertility?click=main_sr/, Oct 2012.

[20] M.P.H Esther Eisenberg, M.D. Polycystic ovary syndrome (pcos) fact sheet. http://womenshealth.gov/publications/our-publications/fact-sheet/polycystic-ovary-syndrome.cfm, March 2010.

[21] Vingerhoets Eugster. The efficacy of psychological interventions for infertile patients: a meta-analysis examining mental health and pregnancy rate. 2009.

[22] Society for Assisted Reproductive Technology. Success rates. http://www.sart.org/SART_Assisted_Reproductive_Technologies/, 2013.

[23] Society for Assisted Reproductive Technology. Success rates. http://www.sart.org/SART_Success_Rates/, 2013.

[24] Center for Disease Control. Heart disease facts. http://www.cdc.gov/heartdisease/facts.htm, March 2013.

[25] Center for Disease Control. Reproductive health: Infertility faq. http://www.cdc.gov/reproductivehealth/infertility/#14, Feb 2013.

[26] Massachusetts General Hospital Center for Women's Mental Health. Infertility, assisted reproduction and mental health. http://www.womensmentalhealth.org/resource/patient-support-services/infertility-assisted-reproduction-and-mental-health/.

[27] Glazebrook Kevin Weber Richard Gittins, John. *Multi-armed Bandit Allocation Indices*. Wiley, 2011.

[28] Godehardt E. Frank-Herrmann et al Gnoth, C. Definition and prevalence of subfertility and infertility. *Human Reproduction*, 20(5):1144–1147, 2005.

[29] M. G. R. et al Hull. Population study of causes, treatment, and outcome of infertility. 1985.

[30] The IVF Informant. Iui cost 2012. http://ivf-info.org/iui-cost-2012, Oct 2012.

[31] John A. Collins Jacky Boivin, Laura Bunting and Karl G. Nygren. International estimates of infertility prevalence and treatment-seeking: potential need and demand for infertility medical care.

[32] H. Leridon J.L.H. Evers B. Lunenfeld J.D.F. Habbema, J. Collins and E.R. teVelde6. Towards less confusing terminology in reproductive medicine: a proposal. *Human Reproduction*, 19(7):1497–1501, 2004.

[33] Hjollund Niels Henrik I et al Jensen, Tina Kold. Does moderate alcohol consumption affect fertility? follow up study among couples planning first pregnancy. *BMJ*, 317(7157):505–510, 1998.

[34] Abha Khetarpal and Satendra Singh. Why cant we classify this inability as disability? 2012.

[35] Daniel Glinoer Kris Poppe, Brigitte Velkeniers. Thyroid disease and female reproduction. *Clin Endocrinol.*, 66(3):309–321, 2007.

[36] Laura Malfere. *Modern Lifestyles: The Newest Fertility Treatments: How Modern Women Can Buy Time*. HyerLink, 1994.

[37] Boerma T Vanderpoel S Stevens GA Mascarenhas MN, Flaxman SR. National, regional, and global trends in infertility prevalence since 1990: A systematic analysis of 277 health surveys. *PLoS Med*, 9(12):1001356, 2012.

[38] National Diabetes Information Clearinghouse (NDIC). National diabetes statistics, 2011. `http://diabetes.niddk.nih.gov/dm/pubs/statistics/#fast`, Dec 2011.

[39] March of Dimes. Miscarriage. `http://www.marchofdimes.com/baby/loss_miscarriage.html`, Oct 2008.

[40] March of Dimes. Stillbirth. `http://www.marchofdimes.com/baby/loss_stillbirth.html`, Oct 2008.

[41] US National Library of Medicine. Infertility. `http://www.ncbi.nlm.nih.gov/pubmedhealth/PMH0002173/`, Feb 2012.

[42] U.S. National Library of Medicine. Polycystic ovary syndrome. `http://www.ncbi.nlm.nih.gov/pubmedhealth/PMH0001408`, Feb 2012.

[43] The Practice Committee of the American Society for Reproductive Medicine. Smoking and infertility: a committee opinion.

[44] Male Infertility Best Practice Policy Committee of the American Urological Association. Infertility: Report on the evaluation of the azoospermic male. Technical report, American Urological Association, 2001.

[45] University of Utah Health Care. Male factor infertility. `http://healthcare.utah.edu/healthlibrary/library/diseases/adult/doc.php?type=85&id=P01484`.

[46] World Health Organization. Infertility definitions and terminology. `http://www.who.int/reproductivehealth/topics/infertility/definitions/en/index.html`, 2013.

[47] Richard Paulson. Patient information: In vitro fertilization (ivf) (beyond the basics). http://www.uptodate.com/contents/in-vitro-fertilization-ivf-beyond-the-basics, April 2013.

[48] Rochelson B Kahn E Peng H, Levitin-Smith M. Umbilical cord stricture and overcoiling are common causes of fetal demise. 2006.

[49] William H. Press. Bandit solutions provide unified ethical models for randomized clinical trials and comparative effectiveness research. *Proceedings of the National Academy of Sciences of the United States of America*, 106(52):22387–22392, 2009.

[50] IMG Health Publications. Chemical pregnancies. http://www.pregnancycorner.com/loss/chemical-pregnancy.html, 2013.

[51] Martin L. Puterman. *Markov Decision Processes*. Wiley, 2005.

[52] Iva Skoch. Should ivf be affordable for all? http://www.thedailybeast.com/newsweek/2010/07/20/should-ivf-be-affordable-for-all.html, July 2010.

[53] H. K. A. et al Snick. The spontaneous pregnancy prognosis in untreated subfertile couples: the walcheren primary care study. 1997.

[54] Banks Catherine M. Sokolowski, John A. *Modeling and simulation in the Medical and Health Sciences*. Wiley, 2011.

[55] Parnell Tracey A. Boyle Phil C. Stanford, Joseph B. Outcome from treatment of infertility with natural procreative technology in an irish general practice. *J Am Board Fam Med*, 21(5):279–295, 2008.

[56] F.A.C.O.G. Steve Eisinger, M.D. Uterine fibroids fact sheet. http://womenshealth.gov/publications/our-publications/fact-sheet/uterine-fibroids.cfm#e, May 2008.

[57] Stanford University. What causes female infertility? http://www.stanford.edu/class/siw198q/websites/reprotech/New%20Ways%20of%20Making%20Babies/Causefem.htm.

[58] YourDays. Everything you need to know about artificial insemination. http://www.yourdays.com/articles/fertility/artificial-insemination/everything-you-need-to-know-about-artificial-insemination.htm.