

The Joy of Learning. Main Ideas, Scaffolding, and Thinking:  
building new concepts by modeling: HOWTO

Josh and Frank Starmer

September 23, 2005

## Abstract

*Walking through life can be fun and interesting, dull and boring or simply frustrating. Our walk is based not on our talents but on our decisions. Our skill in making decisions that add a little spice to each day depends on our ability to bring together facts and concepts, develop a picture of the desired outcome and then make a decision that brings us closer to that outcome. In our experience, making decisions is simplified if one understands what we call "main ideas".*

*We emphasize main ideas because most of us have difficulty remembering unrelated facts. For us, our memory works best when there is a framework, or scaffolding, on which we can hang facts and concepts. The scaffolding provides the links between facts and concepts that help with recalling and manipulating some of the stuff in our memory. The simplest scaffolding is built from main ideas that often reflect an oversimplification of an area, but nevertheless, provides the essential framework for building an understanding of more complex systems.*

*The Internet and efficient search engines allow us to focus our energy on understanding main ideas rather than memorizing all of the facts and concepts associated with a particular area of interest. In the 21st century we can justifiably question cluttering our memory with detailed knowledge when we can retrieve this knowledge from an Internet search and quickly reconstruct the details from our scaffolding built from main ideas. The concept of building our framework by concentrating on main ideas is simply a more efficient way of thinking. Moreover, it is an efficient way of educating oneself. Main ideas feeding the construction of scaffolding, mixed with a good dose of curiosity, provide the substrate for life-long learning and restore the fun of learning. Within this context, education is no longer arbitrary and something we are simply told to do. Instead it is fun because the process of discovering a new main idea, adding to our scaffolding, gaining a new insight and trashing irrelevant knowledge is refreshing.*

*How do we identify what to learn and what we can safely disregard? We find problem-based learning to be the tool for separating the information we need to solve a particular problem from what we can safely ignore. Content-mastery requires we start at the beginning of a book and learn all the stuff between the first and last pages. We can be easily evaluated by testing what we remember. Success with problem-based learning requires we readily identify what we don't know. We identify what we don't know with mental images. If we can paint an unambiguous picture of what we are thinking about then we have no gaps in our understanding. If there are gaps in the picture, then new learning is necessary. If we easily recognize what we don't know then when we start to solve a new problem, we quickly realize the gaps in our tools for solving this particular problem.<sup>1</sup> Our approach here is to provide you with the main ideas or concepts upon which you can acquire enough new knowledge to solve a problem of interest. Our approach to main ideas is based on mastery of three concepts: problem-based learning, Internet-searching and Internet-memory.*

*Here we present main ideas (or central concepts), acquired over the past 35 - 40 years, that have facilitated work in areas ranging from cardiac pharmacology to biostatistics to software engineering. In addition, we present some HOWTO segments that introduce you to tools we find generically useful: octave for modeling, xmgrace for displaying data, cvs for collaboratively developing a document or software. The underlying theme of "main ideas" is the fun of learning derived from the beauty and elegance of simple yet powerful concepts. Simple concepts usually arise from simple questions and simple questions more often arise from children than from adults. Perhaps one of the most challenging activities for us as adults is to rediscover the fun of learning through asking childishly simple questions.*

---

<sup>1</sup>Here our educators have to solve the problem of how to test our problem solving skills instead of our skill in memorizing. This is a topic for another paper.

# Chapter 1

## Why Create Models?

### 1.1 Introduction

Problem-solving is the core of a service-based economy. Problem solving requires thinking and thinking is best accomplished within some frame of reference which we call scaffolding. We have found that some educators are beginning to realize the importance of scaffolding in the process of critical thinking, but we have found very little about identifying individual components of the scaffolding. We believe (actually assert) that main ideas, i.e. the underlying theme of some process or activity, is a piece of scaffolding. Model building is a means for focusing on and identifying “main ideas,” a concept that promotes efficient problem solving and is becoming more and more vital in our world. Thus, before we explain why one would want to create models, let us first explain the significance of main ideas.

### 1.2 Education: The Proper Balance of Memorizing, Understanding and Thinking

For the US, Canada and Europe, the 21st century has seen a shift in our industrial base from a manufacturing-based economy to a service-based economy. Service, of course, is problem solving and problem solving requires thinking. If we find a way to make problem-solving both fun and profitable, then it is highly likely that we’ll enjoy an interesting life.

Our brain is capable of four primary functions: memorizing, where we store facts and concepts, understanding, which is how we take a concept and extend it to a different setting, thinking, where we build new concepts from stored facts and concepts and do-

ing, where we perform some task. There is considerable evidence that storing information results in structural changes in our brains. Thus each fact we memorize creates some sort of structure that we can draw on for future use. Because there is a structural basis of memory, we must realize that unlearning a fact or a reflex requires possibly more energy than that required for the initially storing the fact or developing the reflex.

### 1.3 A digression: the biological basis of memory and learning

Eric Kandel and his laboratory [?, ?, ?] used the sea snail (aplysia) to demonstrate the mechanisms that implement neuronal learning. Exploring the siphon withdrawal reflex<sup>1</sup>, they found that learning was activity-dependent. That is, training (repeated stimuli) that occurs frequently appeared to result in more reliable learning than infrequent stimulation.

Since memory is not magic, then there must be a change within a cell or within a network of cells that reflects the learned or memorized element. Our guess is that short term memory is probably the result of in-

---

<sup>1</sup>The siphon is used to facilitate the snail’s breathing. When aplysia breathes, water is drawn across the gill from the front and exits through the siphon. The siphon is usually outside of the snail’s shell or mantle. However, when gently touched, the snail will withdraw and protect its siphon for a short period of time. If this touch is preceded by an electric shock to the tail, the snail will withdraw its siphon for a longer period of time. The snail will continue to have this exaggerated response for up to a day following the shock, and thus, is an example of short term memory. Multiple shocks given over multiple days cause this exaggerated response to become even more exaggerated and retained for much longer. Even after a week since the electric shocks, the snail will continue to exhibit the exaggerated behavior, and this is an example of long term memory.

tracellular accumulation of calcium<sup>2</sup>, Ca, associated with repeated training episodes. Long term memory probably reflects structural changes, either an amplification or attenuation of expression of a cellular signal receptor or channel.

If short term memory is the result of the accumulation of intracellular Ca, then there must be some means for controlling this accumulation. An action potential<sup>3</sup> of some duration is generated each time a cell is excited by a suprathreshold stimulus. The action potential represents the competition between inward (depolarizing) and outward (repolarizing) currents and its duration reflects the magnitude of the net inward or net outward current (net current = inward - outward). An inward sodium ion, Na<sup>+</sup>, current flowing through sodium channels is usually responsible for initially depolarizing a neuron and generating an action potential. An outward K current is usually responsible for restoring the charge balance that is altered by the influx of Na. Because calcium channels are open at potentials more positive than the sodium channel activation potential, they are open during a large portion of the action potential. The length of time the calcium channels remain open is determined by the duration of the ac-

---

<sup>2</sup>We are assuming calcium, Ca, is the agent responsible for short term memory because it is probably critical for neurotransmitter release required to signal adjacent neurons. More intracellular Ca would trigger an increase in the release of pre-synaptic neurotransmitter which would then activate more post-synaptic receptors, giving a larger post synaptic response.

<sup>3</sup>A brief primer in electrophysiology: *Membrane Potential* is the result of a difference in the relative concentrations of positively and negatively charged particles on opposite sides of a cell's plasma membrane. Cells that have the ability to maintain an a transmembrane gradient in charged ions (and thus generate a membrane potentials) and can rapidly change their membrane potential following a suprathreshold stimulation are called *excitable cells*.

An excitable cell is either in the rest state where the transmembrane potential is -50 to -80 mV or in the excited state where the transmembrane potential can become as large as +40 mV for a few milliseconds. The electrical response to suprathreshold stimulation is called an action potential (see Figure ??) and is caused by the rapid influx of a + charge carrier (either Na or Ca). The restoration of the charge balance is accomplished by a slower efflux of a + charge carrier (K) from the intracellular fluid. However, the charge redistribution alone is insufficient to keep the cell healthy.

The charge carriers must also be redistributed - which is a regulatory process that takes place in the background of cellular activity. Because charge flow during the action potential is down concentration gradients, it is physically impossible to restore the charge carriers without *active transport* up the concentration gradient by actively exchanging ions between the extracellular and intracellular fluids. The Na-K transporter is an example of an active transporter that exchanges extracellular K for intracellular Na.

tion potential, which in turn is determined by the amplitude of the repolarizing potassium current(s). Large K<sup>+</sup> currents will rapidly repolarize the cell while small K<sup>+</sup> currents will prolong the duration of the action potential. Consequently, the calcium influx can be controlled by altering the the availability of potassium channels. Some potassium channels are activated by intracellular second messengers. These receptor-linked channels are sensitive to extracellular neurotransmitters (ACh, serotonin, GABA, dopamine) and provide a way of communicating extracellular events (presence of a neurotransmitter in the vicinity of a membrane-bound receptor) to an intracellular process (generating an action potential).

When the time between training episodes (activation of a channel or activation of a receptor) is greater than the time required to restore the balance of intracellular agents reflecting short term memory, then there will be no accumulation and thus no learning. Its easy to understand why rereading a poem 1/year is less likely to result in memorizing the poem than rereading it 1/hour. If, on the other hand, the time between training episodes is less than the restoration time of the intracellular Ca, then the intracellular [Ca] will increase until a threshold is reached (my guess). The protein expression machinery is probably activated when this threshold is exceeded - resulting in either expressing something new, or amplifying/attenuating the expression of existing cellular component, for example, a receptor, such that the cell will have a permanently increased sensitivity to a circulating neurotransmitter. We speculate that short term memory reflects accumulation of something while long term memory reflects a structural change in the cell or cellular network.

Given this view, its interesting to look at forgetting and unlearning. The less something is used, the more likely it will be forgotten. On the other hand, something that is frequently used will be reinforced. To unlearn a frequently used habit thus requires significant expenditure of effort - either to reverse the structural changes or to disable the structural changes. If someone want's a great research project - we suggest exploring the molecular process of forgetting - to complement Kandel's work on remembering.

This model suggests a main idea: Repetition is essential for most types of learning. We'll simply state that Repetition is the first law of learning!

## 1.4 Trading memorizing for thinking

It is obvious that each minute spent memorizing is one less minute available for thinking. Thus, it is reasonable to ask, “what is the appropriate balance between time spent learning, memorizing and thinking?”

Learning takes place in many forms, we are sure, but for us, learning is easiest when we have a frame of reference, a sort of mental scaffolding or mental picture, on which we can hang new facts and extend our understanding. Main ideas are the structural elements of our mental picture or scaffold. Usually, we start with an over-simplified view of some concept - with just the most primitive framework. We then add detail as necessary that completes enough of our picture so that we can readily visualize the answer to some question (our problem). Its obvious that we will not often need the completed picture in order to answer a question - but only that segment of the picture that is relevant to our immediate question.

Where do we obtain the detailed elements used to build our picture? The internet and commodity computing provides an adjunct to our memory in such a way that our faulty human memory can be replaced, in many respects, by a faultless internet memory. These computer memories, distributed around the world and connected by the internet, create the possibility of an internet-centric work environment where our thinking is augmented by internet access to information resources around the world. If you accept this paradigm, then education should shift its emphasis from traditional learning (with a dominant memorizing component) to problem-solving where thinking dominates.

Problem solving is facilitated by:

- Clear problem statement.
- Subdividing the problem into atomic components.
- Solving (experimentally) components, while building a solution that combines the the components in a way that can be evaluated at each stage of development.
- Utilizing information obtained from as many sources as possible for the solution synthesis process.

The last item indicates that access to information is critical for problem solving and thinking. In our opinion, the ability to use google represents an essential skill for problem solving worthy of association with reading, writing and arithmetic. Google is the 21st century tool that enables us to capitalize on machine memory of distant facts and bring them to bear on our problems of today.

## 1.5 Problems Vs. Disciplines

As we progress in our education, we find ourselves specializing in a specific discipline. Some of us study physics, others study genetics or statistics. These different disciplines have, over the years, been created by a rather arbitrary process. Where does biology stop and chemistry start? Or, where does chemistry stop and physics begin? Thus, it comes as no surprise that real problems do not always fall nicely into a single discipline. In fact, it is the gray areas that fall between disciplines that often offer us the most interesting problems because it is only these problems that allow us to see the big picture. As a result, *multi-disciplinary* is becoming the new buzz word in the scientific community.

Hundreds of years ago there were so called *Renaissance Men*. Men and women that understood the details in many different fields and could successfully contribute significant findings to all of them. Over time, however, each discipline became more and more of an island due to the massive amount of discovered knowledge and the impossibility of one man being able to master it all. Each new concept was built on an increasingly formidable hierarchy of existing theory within a particular discipline.

While the general trend for each discipline to become more and more specialized as it becomes rich with knowledge has not changed, our ability to access the information has. When books and libraries were our only source of information, it was not feasible to have collections of reference books for all disciplines within reach on our desk at the same time. Not only was the possibility of having a library that contained them all incredibly small, but our desks were simply not big and strong enough to hold them. Even if we did have access to the books we would need to become our own renaissance man, there was no means to efficiently search through them.

With the internet and search engines, we solve all three problems Our desks only need to be strong enough to support a computer (and often times, a

laptop computer does not even need a desk), almost all of modern scientific knowledge can be accessed through the internet and we can efficiently search through it all with google. Quick and easy access to information allows us to now focus on the main ideas without having to be bogged down with memorizing all of the details. The details are right at our fingertips when we need them.

This body of work that you are now reading represents information obtained from personal experience, which has allowed us to isolate and identify what we are calling “main ideas” from our own thinking, books and the internet. When this project began, almost all of the text was derived from books and personal experience. However, at this point, the vast majority of the writing is inspired by information found on web pages.

## 1.6 What we observe: signal and noise

### **THE MAIN IDEA - Distilling insights from observations:**

Our sensors detect change comprised of an information bearing signal and noise. We process our perceptions with a variety of variance-reducing tools in order to make the signal more obvious. Signal processing, applied to either time or spatial sequences of values, are used routinely in electrical engineering to identify signals or restore signals. Biological sequences naturally arise within the genome and in proteins. The same tools found useful in extracting signal from noise in electrical signals should be useful in extracting biological messages from these sequences. The trick is to discover how to represent an amino acid or a nucleic acid such that the message becomes obvious.

When I run an experiment, I have two types of models running around in the back of my head. One is a model of the underlying process that I am studying, mechanisms of ion channel blockade. This model is a piece of cake. I make measurements, then fit the measurements to the values predicted by the underlying physical model of ion channel blockade.

But also an implicit model is running around in the back of my head. An implicit model that reflect data manipulations I do in order to salvage results from an experiment, that for whatever reason, is not stable. For example, it is well known that in whole cell voltage clamp studies, the preparation “runs down”

over the course of the study. By this I mean that if you do nothing except make a measurement (peak Na current), every minute for 20 minutes, the results will show a gradual reduction of peak Na current - perhaps as much as 10 - 20 percent, sometimes even larger.

### 1.6.1 How Normalization Can Change Your Model

When an experimentalist have rundown in a preparation, the traditional analysis strategy is to “normalize” the data - i.e. to make a measurement, apply the intervention (superfusing the cell with a drug) and divide or subtract (depending on the situation) the first “control” result from each of the measurements made during the intervention. Then the next intervention would be to wash out the drug by superfusing with a drug-free solution. If there has been little rundown, then the peak Na current after washout will be similar to that before the drug was applied. This, however is rarely the case.

With normalization, we are imposing a model on the data. We are assuming that the changes in our measurements are due only to rundown or the intervention and nothing else. Often this is the case, but there are situations where this is not the case.

Consider the study of lidocaine block of cardiac Na channels. It is well known that the fraction of blocked channels changes with the transmembrane potential of the cell. Hyperpolarized cells experience little block whereas depolarized cells experience significant block. Now we study the voltage dependence of lidocaine block. We make a control measurement, apply the drug, make measure the fraction of blocked channels associated with each pulse of a train of depolarizing pulses. Then we depolarize the holding potential, and repeat this protocol: control pulse, train of pulses.

Now, since we know that lidocaine blockade is dependent on the transmembrane potential, (holding potential), if we divide each current measured during the train of pulses by the control pulse, we correct not only for rundown, but also we abolish the known voltage dependence of lidocaine blockade. Superposition of an explicit and implicit model occurs often, simply because the experimenter is unaware that normalization of data actually carries with it, an alteration of data that is equivalent to adding a component to the model. Unfortunately, this hidden addition to the model goes unnoticed and analyses

can lead to incorrect conclusions due to confounding of effects described by the explicit and implicit model.

## 1.7 Models Reveal Main Ideas

### **THE MAIN IDEA - Problem Solving is Modeling:**

The essence of problem solving is building a mental image of the problem. Often the mental image can be characterized with symbols, a graph or an equation related to some physical concept. The process of developing this mental image is called modeling.

We can develop a model at different levels of complexity. We can decide that we want to reproduce behavior at the 17th decimal point of precision, or we can decide that we are comfortable if we only get the direction of the behavior correct. The decision about the level of precision we are trying to capture with our model is a form of abstracting the problem. When we abstract a problem, we attempt to decide what is relevant and what is irrelevant.

Typically, when we create a model, we start with the simplest, first order, behavior. The goal is to try to get this right without worrying if the time and space scales are correct. This is because if we can not get the first order behavior right, then it is a waste of time to try to get the spatial and temporal scales correct. Thus, at different levels of model building, different levels of detail are relevant. Another reason for starting with a first order approximation is that sometimes that is all you need. If, when you press deeper into the problem, the first order model works, and it continues to work for each successive level of complexity, then we have stumbled on a “main idea”.

Even if we are not so lucky, as we try to characterize the abstractions of multiple instances of our problem, we may begin to see common denominators. This common denominator is a “main idea”, and is the scaffolding around which we can build very complex descriptions of what we observe.

For an example of this latter method of uncovering “main ideas”, consider the problem from cardiology of reentrant cardiac arrhythmias. In normal circumstances, the impulse that initiates cardiac contraction forms as a continuous wave that propagates away from the sinus node. Any continuous wavefront in the heart *cannot* become reentrant simply because it will collide with and extinguish itself. On the other hand,

if a wave breaks and becomes discontinuous<sup>4</sup>, then it is possible for the residual wave fragments to evolve into a spiral wave<sup>5</sup>. Reentrant arrhythmias, rapid uncontrollable reexcitations of the heart, are initiated from wave fragments or discontinuous waves. Therefore, forming a spiral front requires that a front arise in a region with asymmetric excitability where propagation succeeds in some directions and is blocked (or fails) in other directions - <sup>6</sup>. Thus, *all* reentrant arrhythmias can be understood as resulting from wave formation in a region with a spatial asymmetry of cellular excitability<sup>7</sup>. If you can identify the source of the asymmetry, then perhaps its possible to correct it. Since this one concept, asymmetric waves form as a result of propagation in a region with a spatial asymmetry in excitability, enables us to have a general idea about an entire class of phenomena, we will call it a “main idea”.

Modeling is thus an essential step toward identifying main ideas, the recurring themes that we see as we examine different, but related, systems. We will see this when we explore excitable cells in cardiac tissue and transcription switches<sup>8</sup> in the DNA of small organisms.

What follows are some of the main ideas we have developed about building and then testing models. We begin with the mathematics and physics required for model building end with statistics for model evaluation. Along the way, we’ll introduce some of the software issues we have faced as we constructed tools that promoted our development of main ideas.

---

<sup>4</sup>Our friend, Valentin Krinsky, was the first to articulate this

<sup>5</sup>Spirals form from fragments because the ends of the fragment propagate more slowly than the interior segments of the wave. Why? because the ends must excite not only the cells in front of them but also the cells to the side - and, because the cell has a limited charge available to excite adjacent cells, more time is required to transfer this charge to the larger *audience* of adjacent cells

<sup>6</sup>This site in the heart is composed of what are called *pacemaker cells*

<sup>7</sup>Four examples for such an asymmetry are: inexcitable obstacles that the wave collides with, cellular coupling, as described by Maddy Spach, dispersion of refractoriness, or a spatially inhomogeneous distribution of potassium channels.

<sup>8</sup>Both excitable cells and DNA transcription involve switches. Switches are either *on* or *off*. The phase plane of any system with two stable states requires a third, intermediate state that is unstable and possibly oscillatory.

# Chapter 2

## How to create a model

### 2.1 Introduction

**THE MAIN IDEA - Scientific Research:**

The goal in scientific research is to identify the underlying mechanism that is responsible for some phenomena. Great joy and fun is derived from identifying generic mechanisms, mechanisms that are shared by many different phenomena. The generic mechanism is the core concept or theme, while the different presentations of this theme are like different variations on that theme.

What is an example of a core concept and some variations? Frank's research over the past 20 years has focused on reentrant cardiac arrhythmias and the potential role antiarrhythmic drugs might play in amplifying the potential for triggering reentrant arrhythmias. Under normal condition, the membrane potential of a group of pacemaker cells oscillates with a frequency of about 1/min. Each time the membrane potential exceeds a threshold, neighboring cells are excited and a wave of excitation propagates away from the pacemaker region. Because the heart is a closed surface, this wave will eventually collide with itself and thereby is extinguished (due to a property called refractoriness).

A reentrant arrhythmia is one where the excitation wave circulates around the heart without colliding with another wave and therefore is capable of re-exciting the heart. Clearly, a continuous front can never become reentrant. However a discontinuous front can evolve into a reentrant process [?]. The variations of the theme of discontinuous fronts are all the different ways one can make a discontinuous wave: by premature excitation, by collision of a front with an obstacle, by excessive front curvature and by encountering non-uniform refractory states. Each variation has specific detail that is required for the mechanism to successfully function within a specific environment.

**THE MAIN IDEA - Models:**

A model is an abstraction of a real world phenomenon. We can never make a perfect model, but we can build models that are sufficiently accurate that it is difficult to distinguish between them and the real world.

Typically, one begins by creating a model of a specific event or phenomenon. However, over time, one might notice that the model applies to other events or phenomena and can be used to answer questions that are completely unrelated to the original intent, thus, demonstrating the potential for generalizing the model. To be able to demonstrate that a model, as representing some physical mechanism, generalizes to describe processes in many different settings is the greatest thrill possible.

The main challenge in biology is identifying processes, mechanisms and developing an understanding of the minimally complex representation. But before starting to model, we ask, about what features the model must represent. For example, we view a living organism as requiring 6 essential processes: metabolism (converting nutrients to energy sources), translation (translating an electrical signal to motion), signaling (transferring the representation of an event from one place to another), replication (duplicating something) and regulation. If we model an organism, then probably these features must be included in the model. Such models are quite useful, because we can use the model of one entity as a template for investigating and characterizing another. For example, at the level of the nucleus, expression could be considered as replication, signaling could reflect the initiators and terminators of expression, metabolism could reflect the supply of raw materials to the expression system etc.

For all these complexities, though, it seems that linear models are adequate to describe many processes. Not that these processes are inherently linear.






Figure 2.1.1: A Saturating Process: A frequently occurring nonlinearity in biological systems. A dose-response curve of a saturable receptor reaction where  $U$  = unoccupied sites and  $O$  = occupied sites.

Rather most likely, the range over which we can explore them is small, and the processes appears linear.

**THE MAIN IDEA - The Essential Nonlinearity in Biology:**

For nonlinear models (and processes) a primary source (often the dominant source) of the nonlinearity is derived from saturable processes. That is to say, there is a finite concentration of receptors, a finite amount of enzyme or a finite amount of neurotransmitter and all of them can be completely used up or saturated.

In the case of a finite number of receptors, with a huge drug concentration, all receptors are occupied - and increasing the concentration does nothing. Similarly, at very small concentrations, no receptors are occupied - and at intermediate concentrations, some receptors are occupied, leading to the sigmoid shaped dose-response curve: represented by  $b = 1/(1 + l/kD)$  where  $b$  is the fraction of bound receptors,  $l$  is the reverse rate constant,  $k$  is the forward rate constant and  $D$  is the drug concentration. When  $D$  is small,  $b = 0$  and when  $D$  is large,  $b = 1$ . Below we'll derive this relationship exactly.

**THE MAIN IDEA - Why Do Linear Models Work?:**

The main idea in analysis: linear models often work simply because the behavior of the Taylor series expansion of the "real" function is dominated by the linear term. That is why linear models in biological research work so often.

Thus (see below) estimating parameters derived from linear models is an important statistical tool.

We'll derive a simple least squares procedure and hypothesis testing concept that is readily generalized to nonlinear and categorical data models.

**THE MAIN IDEA - Statistics: using our models:**

A main idea in statistics is to characterize a process in such a manner that you can test the sensitivity of the process to some intervention and detect it. Linear statistical models work quite well in biology because (in our opinion) the linear component of a Taylor expansion dominates the behavior of the system.

In addition, the fact that many nonlinearities are derived from saturable processes results in simple linear approximations of the saturable process: 3 lines - one for low concentrations, one for intermediate concentrations, and one for near saturable concentrations. Unless you happen to be operating near the knees of a saturable process, linear models work really well.

One of the main products of statistical theory is that parametric procedures (which assume normally distributed variations) usually give the same answers as their non-parametric sisters. This is due to the central limit theorem. The central limit theorem says that sums of random variates are asymptotically normally distributed. My numerical studies indicate that when you have 7 or 8 terms in the series, then asymptotic normality rules the day. Thus, we never worry about the underlying distribution (well almost never) because we are analyzing sums of random variates (mean, variance etc) which are asymptotically normal (or chi square for sums of squared normal variants).

**THE MAIN IDEA - How many experiment must I do?:**

A major idea in statistical design is that if you really have a strong process you are exploring, then you only need 5 experiments and then you test the direction of the result after the intervention. If the probability of the random (no mechanism) outcome is 0.5 (1/2) then if you obtain the same directional result in 5 consecutive experiments, the likelihood of that happening by chance alone is  $(1/2)^5 = 1/32 = .03 < .05$ , the magic type 1 error threshold.

Model building can be based on algebraic equations or differential equations. When do we use which? The main idea with differential equations can be best viewed by comparing with algebraic equations. Solving algebraic equations results in finding points that satisfy the equations, while solving differential equa-




Figure 2.2.1: Graphical Taylor analysis, the approximation includes only the linear term

tions results in finding functions that satisfy the equations. Here we show how differential equations arise in ordinary problems and how to solve a simple first order linear ODE (ordinary differential equation). With these tools, you can run over any boulder.

Modeling is simply translating a physical process into some equations that describe the physical process. Thus, the idea is to get a mental image of the process to be modeled, then using basic physical and chemical concepts, write the ODEs or PDEs that describe the process.

## 2.2 Taylor Series

Often times you are working with awful functions as part of some analysis. They are impossible to integrate or differentiate or to find the roots of. As with many things, the trick is to find a reasonable approximation that is simple enough to work with. The Taylor Series allows you to do just that. With it you can decompose any complex function,  $f(x)$ , into an infinite series. Often, a truncated series will be used for for the analysis. For example, you might end up using only the first two terms of the series (the linear portion of the approximation). The difference between the value obtained from the infinite series and the truncated series can be called "error" or noise.

The Taylor Series, or Taylor Expansion of any function,  $f(x)$  is defined as:

$$f(x) = f(a) + f'(a)(x - a) + f''(a)\frac{(x - a)^2}{2!} + \dots + f^n(a)\frac{(x - a)^n}{n!} + \text{remainder.} \quad (2.2.1)$$

The first two terms are the most important because this is the linear approximation. It says that for any function,  $f(x)$ , you can create an approximation of it around any point,  $a$ , by looking at the slope of the function at  $f'(a)$  and multiplying it by the distance between  $a$  and another point you want to know something about. It is a straight line approximation, the value of the function at the point,  $x = a$  added to the derivative  $\frac{\Delta f}{\Delta x} \Delta x$ .

### Example 2.2.0.1 (Approximate Exponential)

To demonstrate the how the Taylor series can generate an approximation function for  $f(x)$ , we will define  $f(x)$  as an exponential function and create an approximate function around zero. That is, we will let  $f(x) = e^{-\lambda x}$  and  $a = 0$ . Since,  $f'(x) = -\lambda e^{-\lambda x}$ , the linear portion of the Taylor expansion is:

$$f(x) = f(0) + f'(a)(x - 0) = 1 + (-\lambda)(x).$$

If we want to know  $e^{-\lambda x}$  at  $x = 0.1$ , when  $\lambda = 2$  then

$$e^{-1*2} = 1 + (-2)(0.1) = 1 - 0.2 = 0.8.$$

Compare this approximate answer to the correct answer to 6 decimal places, 0.818731.

Note, if  $-\lambda x$  is positive, then the series will diverge. However, when the exponent is negative, the series converges. That is, the signs of each term in the approximation will alternate. Thus, for some analyses, we can replace  $e^{\lambda x}$  with  $1 + \lambda x$  and continue the analysis. ||

#### THE MAIN IDEA - Why mess with a Taylor Series:

The main idea behind using the Taylor Series is that you can make a linear approximation of some function so that you can then analyze the behavior of the linear function. We usually know how to analyze linear stuff because it is often times well behaved. Nonlinear analysis requires tricks and we are not well versed in nonlinear tricks.

### 2.2.1 Reverse Engineering

Sometimes you are faced with an equation that appears to be a guess by someone. Is there a way to figure out what it approximates? This is particularly true with difference equations when someone is approximating an ordinary differential equation or a partial differential equation and they state a bunch

of difference equations and one looks a bit spooky. So here is a way to reverse engineer what is happening. Typically, the equations involve a function and a couple of points, say  $V_{i-1}, V_i, V_{i+1}$ . Now let's expand the function,  $V$ , around point  $i$  as

$$V_{i-1} = V_i - \Delta x * V'_i + \Delta x^2 V''/2 + \dots$$

$$V_{i+1} = V_i + \Delta x * V'_i + \Delta x^2 V''/2 + \dots$$

Now some preliminaries. Suppose we use  $\Delta V/\Delta x$  to approximate a derivative. We see immediately from above that

$$\frac{V_{i-1} - V_i}{\Delta x} = -\frac{dV_{i-1}}{dx} + \Delta x V''/2 + \dots$$

which shows that we estimate the left derivative with an error that is proportional to the 2nd derivative. To get a sort of unbiased estimator of the derivative - subtract equation 4 from 5 and you see:

$$V_{i+1} - V_{i-1} = 2\Delta x \frac{dV}{dx}$$

or

$$\frac{V_{i+1} - V_{i-1}}{2\Delta x} = \frac{dV}{dx}$$

Now this is a nice estimator - Note that dividing each side by  $2\Delta x$  gives a perfect approximation of the derivative - i.e. the higher order terms disappear!

Now let's say that someone uses the difference equation  $V_{i-1} - 3V_i + 2V_{i+1}/\Delta^2$  and we would like to know what it estimates. So we multiply equation 5 by 2 and add them getting

$$V_{i-1} + 2V_{i+1} = 3V_i + \Delta x \frac{dV}{dx} + 3\Delta x^2 \frac{d^2V}{dx^2}$$

so moving the  $3V_i$  to the left we have

$$\frac{V_{i-1} - 3V_i + 2V_{i+1}}{\Delta x} = \frac{dV}{dx} + 3\Delta x \frac{d^2V}{dx^2}$$

The paper where this appeared stated that the above difference equation approximated the first derivative - and it's obvious that it approximates a bit more than the simple first derivative.

## 2.3 Algebraic Models

Some processes are so simple that they can be described in terms of algebraic equations, either explicitly, or implicitly as the solution to a differential equation. Algebraic equations are usually defined by applying some law of physics like conservation of mass

or conservation of momentum or a time or space dependent equation describing the temporal movement of something. For example this is an explicit algebraic model:

$$\text{age} = x - \text{date of birth},$$

where  $x$  is today's date. An example of an implicit algebraic equation is the description of the time course of binding of drug to a receptor. The dynamics is best characterized by a differential equation (equating changes in the fraction of bound receptors to the difference between rates of forming and unforming bound receptors) which has a simple algebraic solution:

$$b = 1 - e^{-(kD+l)t},$$

where  $b$  is the fraction of bound receptors,  $kD$  is the rate of making bound receptors,  $l$  is the rate of unmaking bound receptors and  $t$  is time.

Algebraic models are usually easy to explore because we can simply generate a sequence of values for the independent variable and plot the resulting values of the model's dependent variable.

## 2.4 Ordinary Differential Equations

Models can be built from words or from equations. We usually start with a word model, or qualitative model, just to get the central concepts organized. But qualitative models are difficult to explore and sooner or later, we find ourselves translating actions in our word models into equations that describe the quantitative results of these actions. Tools such as matlab and octave make quantitative models easy to explore. Simple command line tools, `series` and `tf` also give us a means for exploring algebraic models. `series` generates a sequence of numbers of length `num_terms` from `begin` to `end`. This sequence of numbers can then be piped into `tf`, a tool for evaluating an algebraic expression and these results can be piped into a plot tool. Thus, you can set up an easy pipeline with a shell command like this:

```
shell> series begin end num_terms |
      tf "algebraic_expression" | plot
```

and the moral of the story is that numerical tools enable the model to be used as a simulation of the real phenomena and certain hypotheses can be tested

against it. Word models, on the other hand, can describe a process, but are not so easily converted into computer programs and tested.

Differential equations come in all flavors and sizes. They basically have the form

$$\frac{dy}{dt} = f(y, t),$$

where  $f(y, t)$  can be linear, nonlinear, have constant coefficients or variable coefficients. Often times  $y$  is a function of  $t$ . That is,  $y = y(t)$ . The highest order derivative in the equation determines the *order* of the differential equation. Differential equations with only a single independent variable are called Ordinary Differential Equations (ODEs). Those with more than one independent variable are called Partial Differential Equations (PDEs), due to the fact that the derivatives are partial derivatives. The solution to a differential equation is the unknown function  $y(t)$  that you have the derivative for. While this may seem backward, in nature, we can observe how something changes over time, and thus, we can fit a derivative to this data. From this derivative, we then try to determine the original function.

A differential equation describes changes in one variable relative to another variable, and as such, solutions to differential equations are functions that describe the ups and downs of a function. For example,  $y = \sin(t)$  or  $y = Ae^{-bt}$  where  $y$  is the dependent variable and  $t$  is the independent variable. The differential equation:

$$\frac{dy}{dt} = -by \quad (2.4.1)$$

is an equation that says the change in  $y(t)$  for a certain change in  $t$  is negatively proportional to the value of  $y(t)$ . In other words, when  $y(t)$  is large, the slope of the solution ( $dy/dt$ ) is negative and proportional to  $y(t)$  (the proportionality constant is  $b$ ). As  $y(t)$  becomes smaller, the slope becomes smaller.


The solution to Equation ?? is

$$y(t) = Ae^{-bt}, \quad (2.4.2)$$

where the constant  $A$  is determined by the “initial condition”, the value of  $y(t)$  when  $t = 0$ . If  $y(0) = 1$ , then  $y(t) = e^{-bt}$ ; if  $y(0) = 100$ , then  $y(t) = 100e^{-bt}$ . The solution of the differential equation, produces a “class” of similar solutions, and a particular member of that class is identified by the initial condition.

### Example 2.4.0.1 (Building an ODE)

Consider a simple chemical reaction. We have a substance,  $A$ , that spontaneously converts to  $B$  with a rate,  $k$ , while  $B$  spontaneously converts to  $A$  with a rate,  $l$ . Schematically, we can notate this with the equation:



From this, we can describe the change in  $A$  as the proportion of  $B$  that converts into  $A$  minus the proportion of  $A$  that changes into  $B$ . That is,

$$\frac{dA}{dt} = lB - kA. \quad (2.4.3)$$

If we enforce conservation of mass so that the combined mass of  $A$  and  $B$  is always constant,  $A + B = P_{\max}$ , we can now rewrite the Equation ?? as

$$\frac{dA}{dt} = l(P_{\max} - A) - kA. \quad (2.4.4)$$

Without explicitly finding a solution to Equation ??, we can determine what it will be when it is at equilibrium. That is to say, we can determine what proportion of  $P_{\max}$  needs to be comprised of  $A$  such that the amount of  $B$  converting to  $A$  is the same as the amount of  $A$  converting to  $B$ , or  $kA = l(P_{\max} - A)$ . We do this by setting the slope of  $A$  to zero and solving for  $A$ :

$$\begin{aligned} l(P_{\max} - A) - kA &= 0 \\ lP_{\max} - lA - kA &= \\ lP_{\max} - A(l + k) &= \\ A(l + k) &= lP_{\max} \\ A &= \frac{lP_{\max}}{(l + k)} \\ A &= \frac{P_{\max}}{(1 + \frac{k}{l})}. \end{aligned}$$

Thus, if

$$A = \frac{P_{\max}}{(1 + \frac{k}{l})},$$

then the amounts of  $A$  and  $B$  will not change.

Now that we know what the equilibrium is, it is interesting to look at the general solution to Equation ?? because the equilibrium plays a large role in it. Equation ?? can be solved using various methods. In Example ?? we show how to use an integrating factor to solve for  $A(t)$  and the result is:

$$A(t) = \frac{P_{\max}}{(1 + \frac{k}{l})} - \left[ \frac{P_{\max}}{(1 + \frac{k}{l})} - A(0) \right] e^{-(l+k)t}.$$


Notice what happens as  $t$  gets larger and larger. If we take the limit, we get

$$\lim_{t \rightarrow \infty} A(t) = \frac{P_{\max}}{(1 + \frac{k}{I})}$$

and thus, the system converges on the equilibrium. The exponential term simply causes the difference between the initial condition, the amount of  $A$  at time  $t = 0$ , or  $A(0)$ , the equilibrium to become smaller and smaller as time passes. ||

### Example 2.4.0.2 (Two Components as One)

Now, consider a two component reaction,



This is called a second order reaction because the reaction rate depends on the concentration of two components,  $A$  and  $B$ . However, under certain conditions, it can be treated as a first order reaction, like in Example ???. When the concentration of  $A$  or  $B$  is essentially infinite, and there is a small concentration of the other component, then we have a pseudo first order reaction. Here we will show how this is possible from the differential equation.

We start with a 2nd order equation where the rate of formation is determined by the concentration of  $[A]$  and  $[B]$ ,

$$\frac{dC}{dt} = kAB - IC.$$

We assume that  $A$ s collide with  $B$ s at a rate determined by the temperature of the reaction and that a certain fraction of the collisions will result in making a  $C$ . If the availability of  $A$  is infinite so that its concentration never changes, the rate constant  $k$  can be rewritten as a pseudo rate constant  $k_p = kA$ :

$$\frac{dC}{dt} = k_p B - IC,$$

and this allows us to treat the second order reaction as if it were first order. This assumption, that  $A$  is infinite, is often reasonable when  $A$  is some sort of drug compound and  $B$  is a cellular receptor for this compound. ||

## 2.5 Anatomy of a model

We shall start with the Hodgkin-Huxley equations that describe the excitable process of a giant squid

axon. Although much of what follows is our speculation, we suspect that our rationale for each equation is quite similar to theirs.

The equations that we are about to derive are based on the definitions of the current-voltage relationship for different circuit components. Combining circuit elements alters the total current-voltage relationship, and hence, the behavior of the circuit. Here we use  $V$  to be the potential difference across the circuit element,  $I$  to be the current, the amount of charge,  $q$ , that flows per unit time through the element.  $V$ ,  $I$  and  $q$  are functions of both time,  $t$  and space,  $x$ .

1. Ohm's Law:  $I = gV$ , where  $g$  is the conductance (the reciprocal of  $R$ , resistance) and represents the proportionality constant relating current to the difference in potential across a resistor. This implies that the current through a resistor is linearly proportional to the difference in potential across the resistor (the relationship used to describe current through conducting membrane ion channels).
2. Definition of Capacitance:  $q = CV$ , where  $C$  is the capacitance of the circuit element and is the proportionality constant relating charge with potential. This implies that the charge stored within a capacitor is linearly proportional to the difference in potential across the capacitor.
3. Current is the amount of charge that flows/unit time so taking derivatives of the above, we have:

$$I = \frac{dq}{dt} = C \frac{dV}{dt}.$$

Circuits are constructed by parallel or series combinations of resistors, capacitors and inductors, however, there are few biological analogs of inductors and we will ignore them in this context. The differential equations that describe the behavior of a circuit are derived by applying Kirchoff's conservation laws to the circuit:

1. Kirchoff's Current Law: All of the current that flows into a node (an intersection of 2 or more circuit elements) must be equal to the amount of current that flows out of the node. For example, if a circuit element has one input and two outputs and one amp flows into it, then one amp must be distributed between the two outputs.

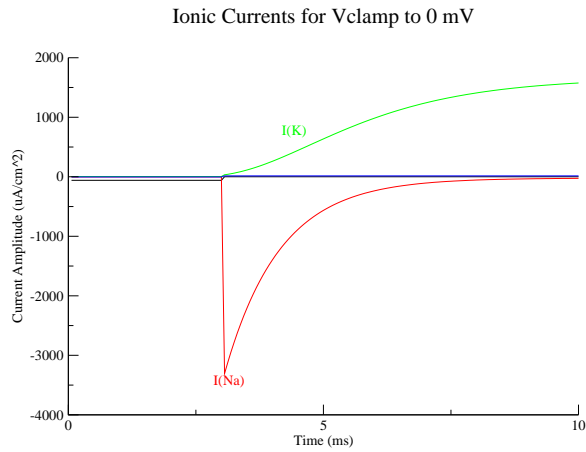


Figure 2.5.1: Computed ionic currents for squid giant axon.

2. Kirchoff's Voltage Law: The sum of the voltage differences measured around a loop of circuit elements must be zero.

Applying these principles to biological systems yields equations that can often characterize a surprisingly large amount of behavior. The fun of modeling is to identify the minimal model required to capture the behavior of a biological process.

In studies of the relationship between current passing across the membrane of a squid giant nerve axon, H-H observed two major currents in response to a step change in the transmembrane potential, an inward Na current that rapidly turned on (activated) and off (inactivated), and a slowly activating outward K current (delayed rectifier) as shown in figure ??.

They incorporated a third current, a leakage current, in order to maintain a balance of current under rest conditions.

Each current was characterized by Ohm's law,  $I = gV$ , but because the ionic currents flowed according to different gradients<sup>1</sup>, the potential,  $V$ , must be related to the reversal potential<sup>2</sup>,  $V_i$ , where  $i$  is simply a

<sup>1</sup>In the presence of both an ionic concentration gradient and an electric field, two currents are possible, one derived from passive diffusion of charge carriers down the concentration gradient and one derived from the attraction of a charge carrier by the electric field.

<sup>2</sup>The reversal potential is the transmembrane potential required to create a current based on charge attraction that exactly balances the diffusive flow of charge carriers down the concentration gradient. For example, consider a higher concentration of Na outside of the cell than inside. The reversal potential required to stop the diffusive current is described by

$$V_{Na} = \frac{RT}{F} \ln \frac{[Na]_o}{[Na]_i},$$

label for the type of current. The total current is thus the sum of two components, the field component,  $g_i V$  and the ion gradient component,  $-g_i V_i$ . Combining the two we can write Ohms law as

$$I_i = g_i(V - V_i).$$

The effect of the gradient current can be seen in Figure ?? where the Na current goes from being negative to positive when the potential is a little over +40 mV. Thus, the reversal potential for Na is +40 mV. For K, the reversal potential is about -80 mV. The sign and the strength of the reversal potential is determined by the different gradients of ions.

H-H considered the membrane as an insulator surrounded on each side by a conductor (extracellular and intracellular fluid). Thus, the membrane acts as a capacitor where the amount of charge that can be stored on the insulating surface is  $q = CV$ . Postulating that the membrane is composed of ion channels that control ion flow between the extracellular and intracellular fluids, the equivalent electrical circuit is the parallel combination of a capacitor, the membrane, and 3 conductances, Na, K and leakage. The current associated with each component of the circuit can be represented by the terms:

	$C \frac{dV}{dt}$	Membrane - capacitive current
$g_{Na} b_{Na}$	$(V - V_{Na})$	Na Channels - channel current
$g_K b_K$	$(V - V_K)$	K Channels - channel current
$+g_L$	$(V - V_L)$	Leakage - channel current

where  $b_{Na}$  and  $b_K$  are the gating terms and represent the fraction of ion channels that are open at a given time. Thus, if all of the Na channels are open, then  $b_{Na} = 1$  and you get full conductance for Na. However, if only half of the channels are open, the conductance is scaled by one half. From Kirchoff's current law, the sum of currents flowing into and out of a node (where circuit element are connected) is zero, we have

$$C \frac{dV}{dt} + g_{Na} b_{Na} (V - V_{Na}) + g_K b_K (V - V_K) + g_L (V - V_L) = 0$$

To extend the model to include propagation, unidirectional movement of ions from cell to cell, we

where  $R$  is the Rydberg constant,  $T$  is absolute temperature,  $F$  is the Faraday constant,  $[Na]_o$  is the concentration of Na outside of the cell and  $[Na]_i$  is the Na concentration inside. The equation is derived by equating the diffusive current with the current created by an electric field. A full treatment of this equation can be found in Appendix ??.




Figure 2.5.2: Flow of diffusive and ionic current within a nerve or cardiac cell. Ionic currents flow down a potential gradient along the radial axis of a nerve or muscle cell, and, as membrane ion channels open, ions can flow down transmembrane concentration and potential gradients. Typically Na and Ca ions flow into the cell while K ions flow out of the cell.

have to add two additional sources of current in this balance, diffusive current into the node and diffusive current out of the node. The current into the node is  $\frac{V(x-\Delta x)-V(x)}{R\Delta x}$ , where  $R$  is the internal resistance per unit length (the reciprocal of  $g$ , conductance). It is easy to imagine that throughout the length of a cell that there would be all sorts of obstacles such as intracellular organs or proteins in the cytosol that would inhibit the free flow of ions through the cell. Since this hindrance is relatively uniform for the different types of ions, we only need one term to account for them. The current out of the node is  $\frac{V(x)-V(x+\Delta x)}{R\Delta x}$ . Because we cannot manufacture charge, then the difference per unit length must equal the current through the membrane.

$$\frac{\frac{V(x-\Delta x)-V(x)}{R\Delta x} - \frac{V(x)-V(x+\Delta x)}{R\Delta x}}{\Delta x} = C \frac{dV}{dt} + g_{Na}b_{Na}(V - V_{Na}) + g_Kb_K(V - V_K) + g_L(V - V_L)$$

Taking the limit as  $\Delta x$  goes to zero then we have the standard nonlinear parabolic partial differential equation with the driving function composed of the individual ionic currents.

$$\nabla^2 V = \frac{\partial V}{\partial t} + \sum I_i \quad (2.5.1)$$

where  $I_1 = I_{Na}$ ,  $I_2 = I_K$  and  $I_3 = I_L$ .

At this point, the two gating variables,  $b_{Na}$  and  $b_K$ , were defined by words, but there was no formula to define their value. These gating parameters separate the H-H model from pure first principles, Ohm's law and conservation of charge. This is where Hodgkin

and Huxley strayed from ordinary science to extraordinary science and this yielded a Nobel prize.

Hodgkin and Huxley's experiments revealed that when they switched the potential across the cell membrane from a polarized value (-60 mV) to a depolarized value, 0 mV, and they poisoned the K charge carriers so that they saw only Na current, the Na current decreased transiently and then returned to zero (or near zero) (the red trace in Figure ??). This feature probably led them to conjecture that there was some sort of dynamic gating process that controlled the flow of ions through the channel. For Na current, they initially suggested two gates, one for activation and one for inactivation. Similarly for potassium, they initially suggested a single activation gate (the green trace in Figure ??).

To test their model they must have plotted observed and expected currents. From these plots they would have observed that the first draft of the model did not fit the initial onset of the activation process for Na or K. For Na current, *three* activation gates resulted in a better fit. The result was that they defined  $b_{Na} = m^3h$ , where  $m$  is the probability that an activation gate opens after the nerve is stimulated, thus  $m^3$  is the probability that three open, and  $h$  is the probability the inactivation gate slowly closes after the nerve is stimulated. Similarly for the potassium current, they found that four gates fit the onset of activation better than a single gate and thus,  $b_K = n^4$ . Substituting terms, our final equation is:

$$C \frac{dV}{dt} + g_{Na}m^3h(V - V_{Na}) + g_Kn^4(V - V_K) + g_L(V - V_L) = 0$$

and the results are shown in Figure ?? where we show the computed action potential and the three gating variables.

This curve fitting exercise is a wonderful example of paying attention to small details. The current voltage relationship of a Na channel is usually measured by holding the cell at some negative (-120 mV) potential and then testing the response with a short duration (5 ms) shift in potential to a test potential. The peaks of each response is plotted and then you think about what the I/V curve is trying to tell you. Shown in ?? are peak Na currents measured in cultured cardiac cells by Gus Grant (+) as a function of the test potential and fits to the I/V curve assuming one (red) and 3 (blue) activation gates. Note that for the initial activation between -60 mV and -45 mV, the red (single) curve overestimates the current

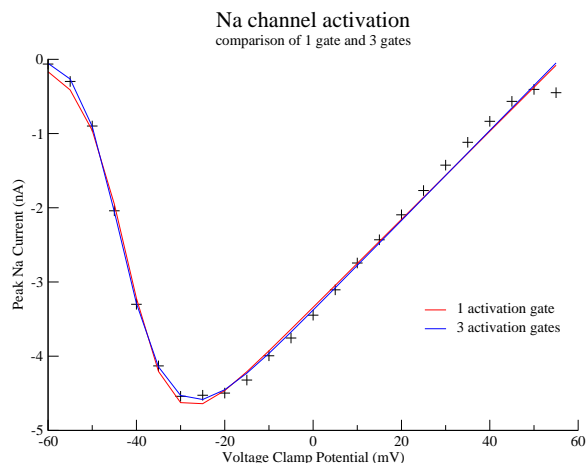


Figure 2.5.3: Comparison of a single and 3 activation gates for cardiac Na channels. Experimentally observed peak currents shown as +, single gate as red and 3 gates as blue

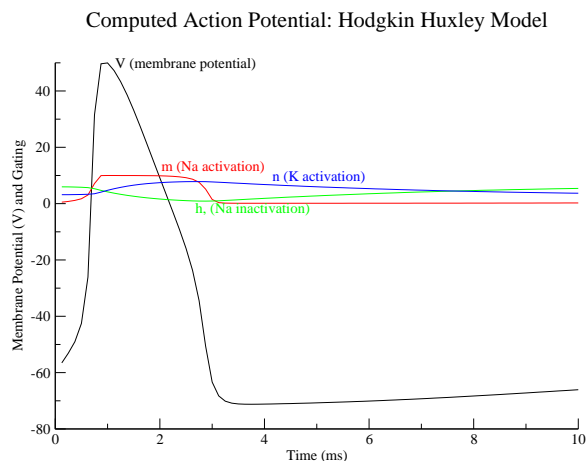


Figure 2.5.4: Computed squid action potential and gating variables.

while the blue curve is right on the money. Similarly, the single activation gate overestimates the peak of the curve. Such consistent overestimation is typically not due to noise, but would suggest that something is missing from the model. In the H-H days of desktop calculators, many (probably including me) would have been happy to get general agreement between observed currents and model currents as shown by the red line, but not Hodgkin and Huxley. They must have realized that there was something not quite right and redid their analysis for a 3 gate process for the Na channel and a 4 gate processes for the K channel.

The first hint that the theory developed by H-H was correct came about 30 years later when the Na

channel was cloned and sequenced by Noma and colleagues in Japan. They observed 4 subunits, each with 6 membrane spanning components and a perfect arrangement for a helical gate. 50 years later, Ray MacKinnon and colleagues managed to crystallize K channels and found four paddles that acted as voltages sensors for the gating process.

## 2.6 Model Approximations and Assumptions

When I run an experiment, I have two types of models running around in the back of my head. One is a model of the underlying process that I am studying, mechanisms of ion channel blockade. This model is a piece of cake. I make measurements, then fit the measurements to the values predicted by the underlying physical model of ion channel blockade.

But also an implicit model is running around in the back of my head. An implicit model that reflect data manipulations I do in order to salvage results from an experiment, that for whatever reason, is not stable. For example, it is well known that in whole cell voltage clamp studies, the preparation "runs down" over the course of the study. By this I mean that if you do nothing except make a measurement (peak Na current), every minute for 20 minutes, the results will show a gradual reduction of peak Na current - perhaps as much as 10 - 20 percent, sometimes even larger.

### 2.6.1 How Normalization Can Change Your Model

When an experimentalist have rundown in a preparation, the traditional analysis strategy is to "normalize" the data - i.e. to make a measurement, apply the intervention (superfusing the cell with a drug) and divide or subtract (depending on the situation) the first "control" result from each of the measurements made during the intervention. Then the next intervention would be to wash out the drug by superfusing with a drug-free solution. If there has been little rundown, then the peak Na current after washout will be similar to that before the drug was applied. This, however is rarely the case.

With normalization, we are imposing a model on the data. We are assuming that the changes in our measurements are due only to rundown or the intervention and nothing else. Often this is the case, but



there are situations where this is not the case.

Consider the study of lidocaine block of cardiac Na channels. It is well known that the fraction of blocked channels changes with the transmembrane potential of the cell. Hyperpolarized cells experience little block whereas depolarized cells experience significant block. Now we study the voltage dependence of lidocaine block. We make a control measurement, apply the drug, make measure the fraction of blocked channels associated with each pulse of a train of depolarizing pulses. Then we depolarize the holding potential, and repeat this protocol: control pulse, train of pulses.

Now, since we know that lidocaine blockade is dependent on the transmembrane potential, (holding potential), if we divide each current measured during the train of pulses by the control pulse, we correct not only for rundown, but also we abolish the known voltage dependence of lidocaine blockade. Superposition of an explicit and implicit model occurs often, simply because the experimenter is unaware that normalization of data actually carries with it, an alteration of data that is equivalent to adding a component to the model. Unfortunately, this hidden addition to the model goes unnoticed and analyses can lead to incorrect conclusions due to confounding of effects described by the explicit and implicit model.

## 2.6.2 Assuming Pseudo-Steady-State

Also known as Quasi-Steady-State, this assumption allows you to model a small portion of an extremely complex system. Simply put, without this assumption many models would not exist and it allows us to work with a system that has both fast and slow reactions. If you are interested in the portion of the model that contains the slow, or rate limiting reactions, you can (sometimes) assume that the fast reactions are in a state of dynamic equilibrium, and the their derivatives are equal to zero, compared to the slow reactions. If you are interested in the portion of the model that has fast reactions, you can assume that the slow portion does not change significantly (and thus, it's derivative is zero) when compared to the fast.

### Example 2.6.2.1 (Focus On Slow Reactions)

For me, in the drug binding business, the channel switches from closed to open and then it can be bound. The binding reaction is slow and the channel transitions are fast. so  $C \xrightleftharpoons[k_2]{k_1} O \xrightleftharpoons[k_4]{k_3} B$  Now

$dO/dt = k_1C + k_4B - (k_2 + k_3)O$  (1,2 are forward and reverse for C-O and 3,4 and forward and reverse for O-B.  $dB/dt = k_3O - k_4B$   $dC/dt = k_2O - k_1C$  but this is fast compared with the others so lets assume that it is always in equilibrium:  $O = k_1/k_2C$  and then using  $1 = C + O + B$ , you know that  $1 = k_2/k_1O + O + B$  or  $O = (1 - B)/(1 + k_2/k_1)$  and now you can substitute into the dB/dt equation for O and solve for B in terms of everything else. ||

### Example 2.6.2.2 (Focus On Fast Reactions)

If we were interested in some intracellular process that required external stimulation, we can often assume that the internal processes are much faster than the external ones due to the fact that the concentrations of the various players in the reactions are much higher internally. ||

## 2.7 Examples of Models

### 2.7.1 Macroscopic/Deterministic Behavior

#### Example 2.7.1.1 (Drug-Receptor Model I)

Let's now look at a few examples of systems that lead to first order differential equations. Consider the process of a neurotransmitter binding to a receptor. Let  $N$  be the concentration of the neurotransmitter, and  $R_o$  be the number of occupied receptors where  $R_{max}$  is the total number of receptors. The number of unoccupied receptors,  $R_u$ , is thus  $R_{max} - R_o$ . Thus, the reaction between neurotransmitter, unbound receptors and bound receptors can be encapsulated with the formula,

$$N + R_u \xrightleftharpoons[k]{l} R_o,$$

where  $k$  is the proportionality constant for binding and  $l$  is the proportionality constant for unbinding. The rate of change of occupied receptors is thus,

$$\frac{dR_o}{dt} = kNR_u - lR_o = kN(R_{max} - R_o) - lR_o. \quad (2.7.1)$$

If we convert Equation ?? to represent the change in the fraction of bound receptors by dividing by  $R_{max}$ , and rearrange the terms a little bit, then we can solve it by using what is called an integration

factor.<sup>3</sup> That is, if we let  $b = \frac{R_o}{R_{\max}}$ , we can rewrite Equation ?? as

$$\frac{db}{dt} = kN(1 - b) - lb. \quad (2.7.2)$$

Rearranging the terms gives us:

$$\frac{db}{dt} = -(kN + l)b + kN, \quad (2.7.3)$$

or

$$\frac{db}{dt} + (kN + l)b = kN. \quad (2.7.4)$$

With Equation ?? in the exact same form as Equation ??, the form required for our general solution, we can easily solve for  $b(t)$ . First, we will determine the integrating factor. That is, from Equations ?? and ??

$$\begin{aligned} \mu(t) &= \exp\left(\int (kN + l)dt\right) \\ &= e^{(kN+l)t}. \end{aligned} \quad (2.7.5)$$

We can now plug the integrating factor, along with bits from Equation ?? into Equation ?? and solve for  $b(t)$ . Thus,

$$\begin{aligned} b(t) &= e^{-(kN+l)t} \left( kN \int e^{(kN+l)t} dt + C \right) \\ &= e^{-(kN+l)t} \frac{kN}{kN + l} e^{(kN+l)t} + C e^{-(kN+l)t} \\ &= \frac{kN}{kN + l} + C e^{-(kN+l)t}. \end{aligned} \quad (2.7.6)$$

Since we know that at time  $t = 0$  that some fraction of receptors are occupied,  $b(0)$ , we can solve for  $C$ . That is,

$$b(0) = \frac{kN}{kN + l} + C$$

and simple rearrangement gives us,

$$C = b(0) - \frac{kN}{kN + l}.$$

Thus, our general solution for  $b(t)$  is

$$b(t) = \frac{kN}{kN + l} + \left( b(0) - \frac{kN}{kN + l} \right) e^{-(kN+l)t}. \quad (2.7.7)$$

||

---

<sup>3</sup>The specific mechanics of this solving an ordinary differential equation using an integration factor is fully described in Section ??.

### Example 2.7.1.2 (Phase Plane Analysis)

If we return to Equation ??, we can determine the asymptotic behavior of the solution without having to solve for it.<sup>4</sup> First, we will determine any points where the system is at equilibrium. That is, determine where the derivative is zero.

$$\begin{aligned} \frac{db}{dt} &\stackrel{\text{set}}{=} 0 \\ kN - (kN + l)b &= 0 \\ b &= \frac{kN}{(kN + l)} \end{aligned}$$

Thus, when  $b(0) = kN/(kN + 1)$ , for all  $t$ , we are at an equilibrium and will not move from it. When  $b(0) < kN/(kN + 1)$ , then the slope for all  $t$  is positive (to see this, try plugging in  $b(0) = kN/2(kN + 1)$ ), and thus, as  $t \rightarrow \infty$ ,  $b(t)$  approaches  $kN/(kN + 1)$  from below. If  $b(0) > kN/(kN + 1)$ , then the slope for all  $t$  will be negative and as  $t \rightarrow \infty$ ,  $b(t)$  approaches the equilibrium from above.

Note that  $b = 1/2$  when  $N = l/k$ . This is called the equilibrium dissociation constant and is the concentration of drug where the fraction of bound receptors is  $1/2$ . Thus, from the kinetics of binding and unbinding, you can directly get a feel for the concentration of drug required to have half the receptors occupied.

||

## 2.7.2 Microscopic/Probabilistic Behavior

### Example 2.7.2.1 (Drug-Receptor Model II)

Example ?? is probably the most useful derivation in biological models. It is also simple with only a few assumptions. If we look at the system more closely we see that binding comes from molecules colliding with receptors and only every once in a while is the collision sufficiently strong that an ‘‘event’’ takes place, producing a complex of ‘‘molecule bound to a receptor.’’ The collisions are due to thermal motion and, up to a certain point, the hotter the solution the more vigorous the collisions and the more binding events. However, beyond that point, the complex can also vibrate and fall apart. Thus, at thermal equilibrium, both binding and unbinding events are constantly happening. In Example ?? we were assuming that the probability of a binding event or an

---

<sup>4</sup>See Section ?? for a complete overview of the general method of phase plane analysis.

unbinding event is constant in time. In this example we will take into account the microscopic view of binding and unbinding and demonstrate that this is a pretty solid assumption to make.

The probability that a receptor is occupied at time  $t$  is  $P_o(t)$  and the probability that a receptor is unoccupied at time  $t$  is  $P_u(t)$ . We also know that  $P_o(t) = 1 - P_u(t)$ . The probability that an unoccupied receptor will become occupied depends on the collision rate which in turn depends on the concentration of hormones, neurotransmitters or whatever molecule is involved in the reaction. We will denote the concentration of the molecule as,  $N$ . The probability that an unoccupied site at time  $t$  will become occupied during the next increment of time,  $\Delta$  is

$$P_{u \rightarrow o} = \lambda N \Delta, \quad (2.7.8)$$

where  $\lambda$  is the proportionality constant for the binding rate per molecule of drug.

For a receptor is bound, there are two possibilities for its state after  $\Delta$ . It can either become unoccupied or remain occupied. The probability that an occupied receptor will become unoccupied during the time interval,  $\Delta$  is a fixed rate,

$$P_{o \rightarrow u} = \mu \Delta. \quad (2.7.9)$$

The probability of an occupied site remaining occupied during  $\Delta$  is simply

$$P_{o \rightarrow o} = 1 - P_{o \rightarrow u} = 1 - \mu \Delta. \quad (2.7.10)$$

To determine the probability that a site will be occupied at time  $t + \Delta$  we need to consider two possibilities. Either the site was empty at time  $t$  and became occupied during  $\Delta$ , or the site was occupied at time  $t$  and it did not become unoccupied during the interval  $\Delta$ . Thus,

$$\begin{aligned} P_o(t + \Delta) &= P_u(t)P_{u \rightarrow o} + P_o(t)P_{o \rightarrow o} \\ &= P_u(t)\lambda N \Delta + P_o(t)(1 - \mu \Delta) \end{aligned} \quad (2.7.11)$$

We can now rearrange terms in Equation ?? and make a difference equation,

$$\begin{aligned} \frac{P_o(t + \Delta) - P_o(t)}{\Delta} &= \lambda N P_u(t) - \mu P_o(t) \\ &= \lambda N (1 - P_o(t)) - \mu P_o(t). \end{aligned} \quad (2.7.12)$$

If we now let  $\Delta$  go to zero, we will end up with a differential equation for the probability that a receptor site will become occupied:

$$\frac{dP_o}{dt} = \lambda N (1 - P_o) - \mu P_o. \quad (2.7.13)$$

If we compare Equation ?? to Equation ?? we notice a striking similarity. Notice that

$$P_o(t) = \frac{R_o(t)}{R_{\max}} = b(t),$$

where  $b(t)$  is the fraction of bound receptors at time  $t$ . Also, both  $\lambda$  and  $\mu$  map to the constants used in Equation ?. Thus, once we take away the assumption that there is a constant event probability over time, we end up with the same general equation. ||

### 2.7.3 A Single Cell

#### Example 2.7.3.1 (Non-Linear and Linear DEs)

Earlier, we briefly explored the role of a cellular action potential plays in short term memory. Cells such as neuronal, cardiac and muscle cells are *excitable*, i.e. when stimulated with a subthreshold stimulus, the cell's electrical potential remains more or less constant. On the other hand, when the cell is stimulated with a suprathreshold stimulus, the cell's electrical potential will change dramatically, and over time return to its rest value.

The action potential is a mechanism for cells communicating with each other. Cellular communication happens when a cell releases a packet of neurotransmitter that binds to a receptor on a nearby cell, or when the cell changes its transmembrane potential and the change is sensed by an adjoining cell. An example of the first method is a nerve cell talking to another nerve cell using synaptic coupling. The latter method could take place in heart cells that are electrically coupled by gap junctions.

What is the minimum complexity of a cell capable of talking (an excitable cell)? In order for it to be useful in signaling, it must be able to have two stable equilibria, rest and excited. By equilibria, we mean points where the derivative,  $dU/dt$ , is zero. By stable, we mean that when you *push* the solution to either side of the equilibrium, the process described by the ode moves the solution back to the equilibrium. In order to have two stable states, the derivative of the current voltage relationship (if we are considering membrane potential as our means of communication) must equal zero at three different conditions, two of these zeros will be stable and one unstable. (See Examples ??, ?? and ?? for illustrations of stable and unstable equilibria.) Therefore a model, driven by a




Figure 2.7.1: Cubic nullcline for Equation ???. Here  $a = 0.25$ .

cubic equation of the form:

$$\frac{dU_B}{dt} = f(U) = U(a - U)(U - 1), \quad (2.7.14)$$

where we assume  $0 < a < 1$ , is required to meet the above conditions. Equilibria exist when this derivative is zero, and thus has equilibria, at  $U = 0$ ,  $U = a$  and  $U = 1$ . A graph of  $f(U)$ , current, is shown in Figure ?? and is called the nullcline. From this, we can graphically explore the behavior of any ode. An equilibrium exists at each point where the nullcline crosses the  $dU/dt$  (current) = 0 axis.

Notice that for  $U < 0$ , the nullcline ( $dU/dt$ ) is positive and for  $0 < U < a$ , the derivative is negative. If the solution is sitting at  $U = 0$  and you push it to the left (say  $U = -0.9$ ) then the value of  $f(U)$  is positive so that  $dU/dt > 0$  and the solution moves back to  $U = 0$ . Similarly, when you push it to the right (say  $U = 0.1$ ), then the value of  $f(U)$  is negative so that  $dU/dt < 0$  and the solution moves to the left, back to  $U = 0$ . Thus,  $U = 0$  is a stable equilibrium. When  $a < U < 1$ , the derivative is again positive and thus  $U = a$  is an unstable equilibrium. Finally, when  $U > 1$ , the derivative is negative making  $U = 1$  a stable equilibrium.

The switching nature of this model can be readily demonstrated. Assume that we are resting at  $U = 0$ . Now as you move  $U$  to the right, the derivative is  $< 0$  so that if we turn the solution loose, it will migrate back to the stable equilibrium at  $U = 0$ . However, if we continue to push so that  $U > a$ , then now  $dU/dt > 0$  and the solution will continue to the right until it reaches the point,  $U = 1$ . The point,  $a$  is called the threshold and with this switch, we have a mechanism for “talking”, switching from a stable rest ( $U = 0$ ) state to a stable excitable ( $U = 1$ ) state.

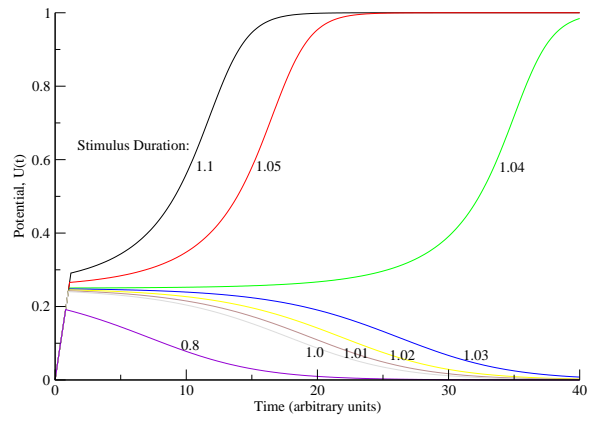


Figure 2.7.2: Switching with near threshold stimulation,  $a = 0.25$

We can add a stimulus function,  $S(t)$  to  $f(U)$  that has magnitude  $m$  by simply creating a function that is equal to  $S$  for a given interval of time. That is:

$$S(t) = \begin{cases} 0, & t \leq 0 \\ m, & 0 < t \leq t_s \\ 0, & t_s < t \end{cases} \quad (2.7.15)$$

where  $t_s$  is the duration of the stimulus. This function could be a reflection of coupling from other cells. For example, a neurotransmitter opening an excitatory channel transiently. Adding  $S(t)$  to  $f(U)$  gives us:

$$\frac{dU_B}{dt} = f(U) + S(t). \quad (2.7.16)$$

Depending on the size of  $m$ , if we initially start at  $f(U) = 0$ , we can switch from the rest state to the excitable state as shown in ??

Here, the stimulus amplitude is 0.25 and the duration of the stimulus is altered. Starting at the stable equilibrium at 0,0, the potential increases linearly until the stimulus value returns to zero. By varying the duration of the stimulus, we can achieve sub-threshold, threshold and suprathreshold values. Shown are the durations of the stimuli. Note that for a duration of 1.4, the value of  $U$  exceeds the threshold (0.25), at the end of the stimulus, and thus, rapidly moves toward the higher stable equilibrium at 1. As the stimulus duration is reduced, the transition time to the equilibrium at  $U = 1$  is progressively longer until the duration is 1.02. Now, the value of  $U$  at the end of the stimulus is no longer suprathreshold and the potential decays back to the stable equilibrium at  $U = 0$ .

Now we must determine how the switch can return to a rest state when it is in an excitable state. Physically, charge is removed from the cell until it crosses the threshold and then the cell takes over, lowering its charge until the lower equilibrium has been reached. As a first attempt at modeling this removal of charge we could define a variable  $V$  such that:

$$\frac{dV}{dt} = aU. \quad (2.7.17)$$

The problem, however, with this definition is that  $V$  will not continue to grow when  $U$  hits the upper equilibrium point. As a result, the system will never return to the rest state. Thus, we must add a second term that will allow  $V$  to grow once  $U$  is at the equilibrium point. In this case we will add a term that causes exponential decay in the charge:

$$\frac{dV}{dt} = aU - bV, \quad (2.7.18)$$

where  $a < b$ . From this equation we can see that as  $U$  increases, so does  $aU$  and thus, so does  $V$ . Once  $U$  plateaus at the higher equilibrium,  $V$  will continue to increase, at slower and slower rates because  $bV$  will continue to subtract from  $aU$  larger and larger amounts, until  $aU < bV$ . Once  $bV$  dominates, Equation ?? will become negative and  $V$  will decrease.

With Equation ?? for  $V$ , our equation defining the entire cell becomes

$$\frac{dU_B}{dt} = f(U) - V + S(t). \quad (2.7.19)$$

By subtracting  $V$  from  $f(U)$  the switch becomes monostable by removing two of the equilibria. We can determine the location of the remaining equilibrium is determined by examining when the derivatives for both equations are zero.

Since  $V$  only plays a significant role after the initial stimulation, we can omit  $S(t)$  in the following derivations. First, we will solve  $\frac{dU_M}{dt} = 0$ :

$$\begin{aligned} \frac{dU_M}{dt} &\stackrel{\text{set}}{=} 0 \\ f(U) - V &= 0 \\ V &= f(U). \end{aligned}$$

Now we will solve  $dV/dt$ :

$$\begin{aligned} \frac{dV}{dt} &\stackrel{\text{set}}{=} 0 \\ (aU - bV) &= 0 \\ bV &= aU \\ V &= aU/b. \end{aligned}$$




Figure 2.7.3: Cubic and linear nullcline for the Fitzhugh Nagumo cell model. When  $U(t = 0) \approx 0.25$ , the potential collapses to the lower equilibrium.

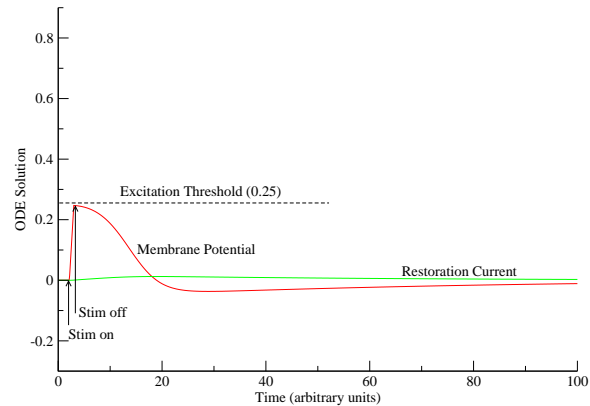


Figure 2.7.4: A response to sub-threshold stimulation at  $t = 2$ . The cell potential,  $U$ , collapses after the end of the stimulation pulse because the phase point did not cross the threshold (0.25) marking the unstable equilibrium ?? marking the transition from  $dU/dt \leq 0$  to  $dU/dt \geq 0$ .

The result is the straight line in Figure ?? with a slope of  $a/b$ . The intersection gives us a single equilibrium.

In biology, the dynamics of moving from a rest state to an excitable state is fast because Na channels open quickly. The recovery, however, is slow because K channels open slowly. We can incorporate this into our model by including a scaling factor into Equation ??:

$$\frac{dV}{dt} = \epsilon(aU - bV). \quad (2.7.20)$$

Keep in mind that  $\epsilon$  does nothing to alter the equilibrium point since it simply divides out when solving for it.

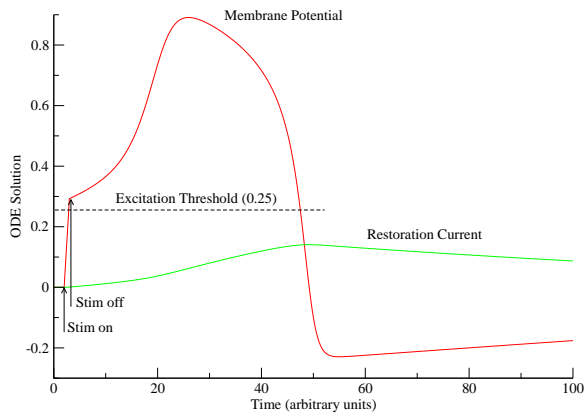


Figure 2.7.5: A response to supra-threshold stimulation  $t = 2$ . In this case, the potential,  $U$ , exceeds the threshold crossing into the region where  $dU/dt > 0$ , and thus accelerates toward the equilibrium at  $U = 1$ . The peak of the action potential never reaches the point where  $U = 1$ , due to the cooling effect reflecting the slow parameter,  $V$ .

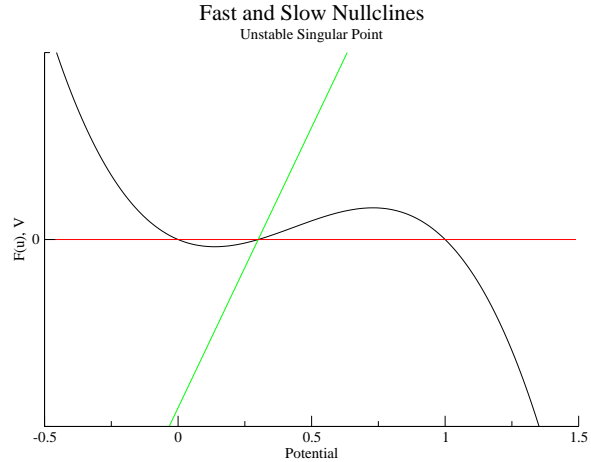


Figure 2.7.6: Shifted equilibrium to the unstable point

The switch defined by Equation ??,  $\frac{dU_B}{dt}$  is bistable with stable equilibria at  $U = 0$  and  $U = 1$ . What we would like to do is incorporate a variable that will remove charge from the cell during the excitable stage until the threshold is crossed and the cell and reset itself to the rest state. ||

Now - move the slow function linear nullcline to the right where the equilibrium is unstable. Now the FHN system behaves as an oscillator as shown in the figure. Same exact model - only a shift in the intersection of the two nullclines ( $f(U) = 0, V = U / \gamma$ )

The result is spontaneous oscillation, because the nullclines intersect at a singular point that is unstable.

## 2.8 Taylor series and identifying generic properties

Now let's use the Taylor series and an arbitrary ordinary differential equation and explore some potentially interesting behavior. First there are two classes (at least) of model builders. Class one is interested in building a full model of some process that is as realistic as possible. Class two is interested in building a minimal model, one that captures essential behavior and upon which, one can add more and more realism and ask: How does this alter the behavior of the minimal model?

We start with the simplest ordinary differential

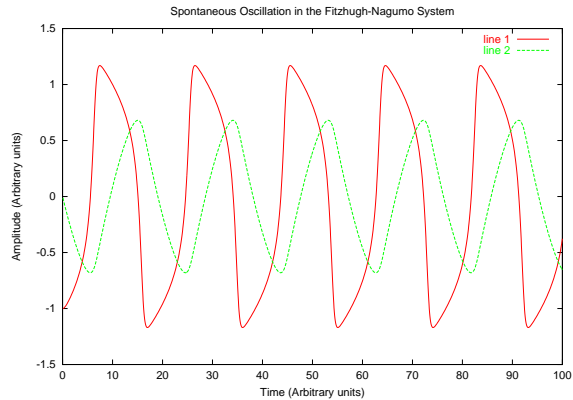


Figure 2.7.7: Oscillation of the FHN system due to the unstable equilibrium

equation:

$$\frac{du}{dt} = u' = f(u). \quad (2.8.1)$$

and ask the question - what is the behavior of this equation as we increase the complexity of  $f(u)$ ? The Taylor series is a way to methodically add complexity (by adding successive terms) in order to more realistically represent a characterization of some unknown function,  $f(u)$ . Starting with the constant term, we can analyze the properties of the ODE and get some ideas about how adequately it represents some process of interest.

So we start with

$$f(u) = f(u_0) + f'(u_0)(u - u_0) + \frac{f''}{2}(u_0)(u - u_0)^2 + \dots$$

The values of  $f(u_0)$  and its derivatives are simply constants so that the Taylor series is simply a power series in  $(u - u_0)$ . So lets rewrite the series as

$$f(u) = a_0 + a_1(u - u_0) + a_2(u - u_0)^2 + a_3(u - u_0)^3 + \dots$$

and start our analysis. For convenience, we will set  $u_0$  to zero.

The properties of  $u' = a_0$  are not interesting. The solutions are lines of varying slope, where the slope is determined by the value of  $a_0$ .

Including the first two terms makes the solution space a bit more interesting:

$$\frac{du}{dt} = u' = a_0 + a_1u \quad (2.8.2)$$

This has a single equilibrium where  $u' = 0$  and the equilibrium is located at  $u = -\frac{a_0}{a_1}$ . Moreover, the equilibrium is unstable<sup>5</sup> if  $a_1 < 0$  as shown in Figure ???. When there is a disturbance that moves the phase point,  $(f(u), u)$  to the left, then we see that  $u' < 0$  which pushes the phase point away from the equilibrium. Similarly when the disturbance moves the phase point to the right,  $u' > 0$  which pushes the phase point to the right, again away from the equilibrium.

Now, we add the quadratic term and have

$$\frac{du}{dt} = u' = a_0 + a_1u + a_2u^2 \quad (2.8.3)$$

We assume that all the constants are such that there are 2 intersections with the  $u' = 0$  line in the  $(u', u)$

<sup>5</sup>See Section ?? for a full explanation of the terms *stable* and *unstable*.




Figure 2.8.1: Linear nullcline for Equation ???. The root is unstable.




Figure 2.8.2: Quadratic nullcline for Equation ???. The left root is unstable, the right root is stable

plane. If there were no intersections, then again, the behavior is not interesting. So these two intersections represent 2 equilibria, one stable and one unstable. See Figure ??.

Next we add the cubic term and have

$$\frac{du}{dt} = u' = a_0 + a_1u + a_2u^2 + a_3u^3 \quad (2.8.4)$$

Now, in the  $(u', u)$  plane, we have constants,  $a_i$  such that there are 3 intersections with  $u' = 0$ . and depending on the values of the  $a_i$ , we either have two stable and one unstable equilibrium or we have two unstable and one stable equilibrium. Figure ?? displays the cubic where we have two stable and one unstable equilibrium.

Now what is interesting about this from a biological modeling perspective? Many biological processes behave like switches. A neuron is either in the rest state or the excited state. A cardiac cell is either in




Figure 2.8.3: Cubic nullcline for Equation ???. The left root is unstable, the right root is stable

the rest state or the excited state - and translate to muscle, the muscle is either resting or contracted. We can even look at transcription. Either the gene being expressed or not.

All of these processes have in common, switch-like behavior. From a modelling perspective, it means that the minimal complex model for describing a switch requires a cubic function on the right hand side of the ODE which means that only nonlinear systems can represent switching behavior. Also, the middle, unstable equilibrium, represents a threshold. So all switches must have a threshold, and we should be able to design experiments to reveal the threshold. Now, a distraction. If there is diffusive coupling between switches and all are initially in the same state, then as one switch is forced to change states, the switches to the left and right can potentially be induced to switch (if the diffusive element forces the local value of  $u$  to exceed the switching threshold) and the result will be a propagating wave.

It is exciting to see the verification of a theoretical argument (above Taylor expansion of an arbitrary function) in real biological systems. In figure ?? we see the current voltage relationship measured in an isolated rabbit cardiac atrial cell. Using the voltage clamp procedure, the potential was gradually increased from negative to positive and the current associated with each potential was recorded. The resultant  $i/v$  is a quasi-steady state and does not accurately reflect the dynamics of a cardiac (or nerve) cell. Nevertheless, the cubic nature is clearly seen (due to calcium channels).

Now the fun part of modeling is to link the cubic function to some real mechanism. In the case of car-

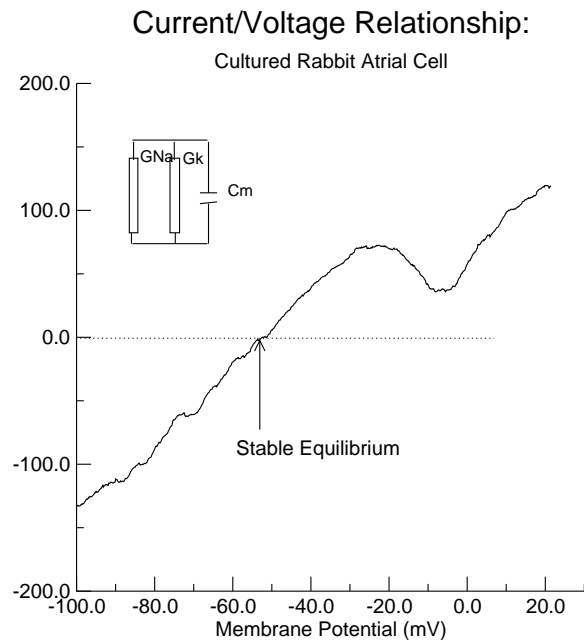


Figure 2.8.4: Current voltage relationship obtained from voltage clamp studies of cultured rabbit cardiac atrial cells. Note the cubic-like behavior

diac and neuronal cells, the cubic function represents the instantaneous current-voltage relationship of the cell. We are unsure what the cubic function represents in a gene expression system.

## 2.9 Analytical Methods for Solving First Order ODEs

The order of a differential equation is determined by the highest derivative that appears in the equation. A first order differential equation has only the first derivative of the function of interest. These equations take the (general) form:

$$\frac{dy}{dx} + P(x)y = Q(x), \quad (2.9.1)$$

where  $P(x)$  and  $Q(x)$  are known functions and  $y$  is an unknown function of  $x$ .

### 2.9.1 Graphical Solutions: Phase Plane Analysis

Oftentimes we can gather a good deal of qualitative information about a solution to a differential equation without going through the trouble of finding an



analytic or numerical solution. Instead, we can simply look for equilibrium points, or points where the derivative is zero, and determine whether the function moves toward or away from these points as time passes giving us the asymptotic behavior of the function without having to solve for it.

### Example 2.9.1.1 (One Stable Point)

Consider the differential equation:

$$\frac{dy}{dt} = -2y - 4. \quad (2.9.2)$$

When  $y(t) = -4/2 = -2$ ,  $dy/dt = 0$ . Thus, if  $y(0) = -2$ , then for any value of  $t$ ,  $y(t) = -2$ , since the derivative will always be zero. Thus  $y(t) = -2$  is called an equilibrium since it will not change. However if  $y(0) > -2$ , then the derivative will be negative, and thus, as  $t$  grows larger and larger,  $y(t)$  will converge to  $-2$ . This is easily seen by simply plugging in different values for  $y$  that are greater than  $-2$ . For example, if  $y = 0$ , then  $dy/dt = -4$ . If  $y = 100$ , then  $dy/dt = -196$ . Likewise, if  $y(0) < -2$ , then the derivative will be positive for all values of  $t$  and  $y(t)$  will approach  $-2$  from below as  $t$  goes to infinity. Since the line  $y(t) = -2$  is approached from above when  $y(t) > -2$  and below when  $y(t) < -2$ , it is called a *node* or a *stable state*. This is because small perturbations to the system at this point will only lead back to it. That is, if the system is at  $y(t) = -2$  and some outside force knocks it to  $y(t) = -1.98$  or  $y(t) = -2.02$ , it will asymptotically return to  $y(t) = -2$ .

Figure ?? shows an actual plot of the phase lines, or slopes for various values of  $t$  and  $y(t)$ . Due to the fact that there are no free instances of  $t$  in Equation ??, the slopes are the same for each value of  $t$ . In this illustration, it is easy to see the stable point  $y(t) = -2$  and how the slope of any point above or below this line points toward it. Figure ?? demonstrates how that regardless of the initial condition, as  $t$  gets larger, the solution will converge on the stable equilibrium. Figure ?? shows how a nullcline graph represents the same information. ||

### Example 2.9.1.2 (1 Stable and 2 Unstable)

Consider the cubic differential equation:

$$\frac{dy}{dt} = (y - 1)(y - 2)(y - 3), \quad (2.9.3)$$

where  $dy/dt = 0$  when  $y(t) = 1$ ,  $y(t) = 2$  and  $y(t) = 3$ . By plugging in different values for  $y$ , we can




Figure 2.9.1: A Phase Plot for Equation ???. There is a single, stable equilibrium at  $y(t) = -2$ .




Figure 2.9.2: A Phase plane for Equation ??, with the solutions for the initial conditions,  $t = 0$ ,  $y(t) = 1$  and  $t = 0$ ,  $y(t) = -3$ . Notice how, regardless of whether or not the initial condition puts  $y(t)$  above or below the stable equilibrium, as  $t$  grows, they both converge to it.




Figure 2.9.3: A nullcline for Equation ?? . When  $t < -2$ , the derivative is positive. When  $t = -2$ , the derivative is zero, and thus, this is an equilibrium point. For  $t > -2$ , the derivative is negative. Since the derivative is positive on the left and negative on the right,  $y(t) = -2$  is a stable equilibrium




Figure 2.9.4: A Phase plane for Equation ??, equilibria at  $y(t) = 1$ ,  $y(t) = 2$  and  $y(t) = 3$ . The equilibrium at  $y(t) = 2$  is stable since the slopes immediately above and below it converge to it.

determine the slope at different points. In this case we end up with one stable point,  $y(t) = 2$  and two *unstable* points,  $y(t) = 1$  and  $y(t) = 3$ . By unstable we mean that if the system is at  $y(t) = 1$  or  $3$ , and is perturbed slightly, it will not return to its original state. Instead, it will either move toward  $y(t) = 2$  or  $\pm\infty$ . This is illustrated in Figure ?? . Figure ? shows the equivalent phase information contained in a plot of the nullcline. When the initial conditions are known, specific solutions can be plotted and this is shown in Figure ?? . ||

### Example 2.9.1.3 (2 Stable and 1 Unstable)

Consider the cubic differential equation:

$$\frac{dy}{dt} = (y - 1)(y - 2)(3 - y), \quad (2.9.4)$$

where  $dy/dt = 0$  when  $y(t) = 1$ ,  $y(t) = 2$  and  $y(t) = 3$ . Again, by plugging in different values for  $y$ , we can determine the slope at different points. In this case, we end up with two stable points and one unstable point. This is illustrated in Figure ?? . ||




Figure 2.9.5: A nullcline for Equation ?? . There is one stable equilibrium at  $t = 2$ , where the derivative to the left is positive and the derivative on the right is negative. The other two equilibria are unstable.

## 2.9.2 Separation of Variables

How do we solve Equation ???. That is, how do we determine what Equation ?? is the derivative of? First, we note that we can separate the two variables,  $y$  and  $t$ , by multiplication. That is,

$$\frac{dy}{y} = -bdt.$$

Integrating both sides produces

$$\ln(y) = -bt + K, \quad (2.9.5)$$

where  $K$  is the combination of the two integration constants. Using each side as an exponent, we have

$$\begin{aligned} e^{\ln(y)} &= e^{-bt+K} \\ y &= e^K e^{-bt} \\ &= C e^{-bt}, \end{aligned}$$

where  $C = e^K$ .

This method can be used to solve both linear as well as nonlinear ordinary differential equations. Example ?? gives an example of a solution to a quadratic nonlinear ODE and Appendix ?? shows how to use separation of variables to solve a cubic nonlinear ODE.




Figure 2.9.6: A Phase plane for Equation ??, with solutions with the initial conditions set to  $y(0) = 3.1$ ,  $y(0) = 3$ ,  $y(0) = 2.5$ ,  $y(0) = 1.5$  and  $y(0) = 0.9$ . Notice how even though  $y(t) = 3$  is an unstable equilibrium, and thus, solutions will not converge on it, if that is where your solution begins, it will not deviate from it without some other force acting on it.




Figure 2.9.7: A Phase Plane for Equation ?. There are two stable equilibria at  $y = 1$  and  $y = 3$  and an unstable equilibrium at  $y = 2$ .

### Example 2.9.2.1 (Quadratic ODE)

We will begin with a quadratic ODE that is often used to model population growth (birth and immigration) and decay (death and emigration).

$$\frac{dy}{dt} = y(k - y) \quad (2.9.6)$$

Equation ?? can be solved using the method of separation of variables. We begin by separating  $y$  from  $t$  by multiplication. That is,

$$\begin{aligned} \frac{dy}{dt} &= y(k - y) \\ \frac{dy}{y(k - y)} &= dt \\ \int \frac{dy}{y(k - y)} &= \int dt \end{aligned} \quad (2.9.7)$$

The integral on the right-hand side can be easily solved once it is broken down into simpler components. This can be done using the method of partial

fraction decomposition. That is,

$$\begin{aligned}\frac{1}{y(k-y)} &= \frac{A}{y} + \frac{B}{k-y} \\ 1 &= \frac{Ay(k-y)}{y} + \frac{By(k-y)}{k-y} \\ 1 &= A(k-y) + By \\ 1 &= Ak + y(B-A)\end{aligned}$$

and  $A = \frac{1}{k}$  and  $B = \frac{1}{k}$ , thus,

$$\frac{1}{y(k-y)} = \frac{1}{ky} + \frac{1}{k(k-y)} \quad (2.9.8)$$

$$\frac{1}{y(k-y)} = \frac{1}{k} \left( \frac{1}{y} + \frac{1}{k-y} \right) \quad (2.9.9)$$

Substituting Equation ?? for the fraction in on the left side of Equation ?? gives us the following:

$$\begin{aligned}\frac{1}{k} \int \left( \frac{1}{y} + \frac{1}{k-y} \right) dy &= t + C \\ \frac{1}{k} (\log |y| - \log |k-y|) &= t + c \\ \log \left( \frac{|y|}{|k-y|} \right) &= kt + C \\ \frac{|y|}{|k-y|} &= e^{kt+C} = e^{kt} e^C = C e^{kt} \\ \frac{|k-y|}{|y|} &= C e^{-kt} \\ \frac{k}{y} - 1 &= C e^{-kt} \\ \frac{k}{y} &= 1 + C e^{-kt} \\ y &= \frac{k}{1 + C e^{-kt}}\end{aligned}$$

||

### 2.9.3 Integrating Factors

Using integrating factors is a useful trick to use when you can not separate the different variables or, if you are able to separate the variables, the integration is too difficult. Here we will present an example of how integrating factors are used and then give a general formula for this method.

The simplest ODEs are linear with constant coefficients:

$$\frac{dy}{dt} = f(y, t) = -ay - b \quad (2.9.10)$$

where  $a$  is a constant. You may have guessed that this is similar to Equation ??, with a similar solution, the exponential function (Equation ??). The only difference here is that there is an added constant. Thus, just as you would expect, the solution is exponential plus a constant. Solving this equation, however, uses a trick, an integrating factor, and in this case the integrating factor is  $e^{at}$ . The main idea is to multiply the equation by a well chosen integrating factor that makes the integration simple. The task of choosing a good integrating factor can be boiled down to following a standard formula, so, overall, finding the solution is not too hard. Thus, starting from Equation ??, we can move everything to one side,

$$\frac{dy}{dt} + ay + b = 0 \quad (2.9.11)$$

and multiply through by  $e^{at}$ :

$$e^{at} \frac{dy}{dt} + aye^{at} + be^{at} = \frac{d}{dt}[ye^{at}] + be^{at} = 0.$$

We can now integrate both sides of the equation and get:

$$ye^{at} + \frac{b}{a}e^{at} + K = 0$$

where  $K$  is an arbitrary integration constant that is determined by the initial conditions.

To finally solve the equation we multiply through by  $e^{-at}$  giving us:

$$y + \frac{b}{a} + Ke^{-at} = 0,$$

or

$$y = -Ke^{-at} - \frac{b}{a}. \quad (2.9.12)$$

To solve for  $K$ , we let  $t = 0$  and thus,

$$y(0) = -K - \frac{b}{a},$$

or

$$K = y(0) + \frac{b}{a},$$

and the complete solution is written as

$$y = -(y(0) + \frac{b}{a})e^{-at} - \frac{b}{a}. \quad (2.9.13)$$

To verify our solution is correct we can take the derivative of Equation ??, plug it into Equation ?? and make sure that everything cancels out. That is,

$$\begin{aligned}\frac{dy}{dt} &= a(y(0) + \frac{b}{a})e^{-at} \\ &= (ay(0) + b)e^{-at}\end{aligned} \quad (2.9.14)$$

and from Equation ??

$$\begin{aligned} ay &= -\frac{dy}{dt} - b \\ &= -(ay(0) + b)e^{-at} - b. \end{aligned} \quad (2.9.15)$$

Plugging the results of Equations ?? and ?? into Equation ??, we have

$$\begin{aligned} \frac{dy}{dt} + ay + b &= (ay(0) + b)e^{-at} - (ay(0) - b)e^{-at} - b + b \\ &= 0, \end{aligned}$$

which is exactly what it should reduce to.

In general, given the equation

$$\frac{dy}{dx} + P(x)y = Q(x), \quad (2.9.16)$$

the solution for  $y(x)$  can be found with the formula

$$y(x) = [\mu(x)]^{-1} \left( \int \mu(x)Q(x)dx + C \right), \quad (2.9.17)$$

where  $\mu(x)$  is the integrating factor and

$$\mu(x) = \exp \left( \int P(x)dx \right). \quad (2.9.18)$$

### Example 2.9.3.1

In Example ?? we created the ODE:

$$\frac{dA}{dt} = l(P_{\max} - A) - kA,$$

which can easily be solved using this method. Following our recipe, we have:

$$\begin{aligned} \frac{dA}{dt} + A(l+k) - lP_{\max} &= 0 \\ e^{(l+k)t} \frac{dA}{dt} + A(l+k)e^{(l+k)t} - lP_{\max}e^{(l+k)t} &= 0 \\ \frac{d}{dt} A e^{(l+k)t} - lP_{\max}e^{(l+k)t} &= 0, \end{aligned}$$

and after integrating both sides with respect to  $t$ , we have:

$$\begin{aligned} A e^{(l+k)t} - \frac{lP_{\max}}{(l+k)} e^{(l+k)t} + K &= 0 \\ A - \frac{lP_{\max}}{(l+k)} + K e^{-(l+k)t} &= 0 \end{aligned}$$

$$A = \frac{lP_{\max}}{(l+k)} - K e^{-(l+k)t}.$$

We can now solve for  $K$  by setting  $t = 0$ :

$$\begin{aligned} A(0) &= \frac{lP_{\max}}{(l+k)} - K \\ K &= \frac{lP_{\max}}{(l+k)} - A(0), \end{aligned}$$

making our general solution:

$$A(t) = \frac{lP_{\max}}{(l+k)} - \left[ \frac{lP_{\max}}{(l+k)} - A(0) \right] e^{-(l+k)t},$$

or

$$A(t) = \frac{P_{\max}}{(1 + \frac{k}{l})} - \left[ \frac{P_{\max}}{(1 + \frac{k}{l})} - A(0) \right] e^{-(l+k)t}.$$

||

## 2.10 Analytical Methods for Solving Second Order ODEs

Second order differential equations involve a second derivative of the function of interest. These equations take the (general) form:

$$P(x) \frac{d^2y}{dx^2} + Q(x) \frac{dy}{dx} + R(x)y = G(x),$$

where  $P(x)$ ,  $Q(x)$  and  $G(x)$  are all known continuous functions and  $y$  is a function of  $x$ . If  $G(x) = 0$ , then Equation ?? is called *homogeneous*. If  $G(x) \neq 0$ , then it is called *nonhomogeneous*.

Generally, solutions to second order ODEs are hard to come by unless  $P(x)$ ,  $Q(x)$  and  $R(x)$  are all constant functions.

### 2.10.1 Homogeneous ODEs with Constant Coefficients

A homogeneous differential equation with constant coefficients has the form:

$$a \frac{d^2y}{dx^2} + b \frac{dy}{dx} + cy = 0, \quad (2.10.1)$$

where  $a$ ,  $b$  and  $c$  are constants and  $a \neq 0$ . Solutions for Equation ?? can be found by first solving for  $r$  in the equation

$$ar^2 + br + c = 0. \quad (2.10.2)$$

Equation ?? is called the *characteristic equation* or *auxiliary equation* and it is obtained from Equation

?? by replacing  $y''$  with  $r^2$ ,  $y'$  with  $r$  and  $y$  by 1. Solutions for  $r$  can be found using the quadratic formula:

$$r_1 = \frac{-b + \sqrt{b^2 - 4ac}}{2a}$$

$$r_2 = \frac{-b - \sqrt{b^2 - 4ac}}{2a}.$$

Solutions for Equation ?? can then be created from the solutions for  $r$ . Depending on the particular solutions for  $r$ , there are three possible forms for the solution to Equation ??:

1. If  $r_1 \neq r_2$  and both are real valued solutions, that is to say,  $b^2 - 4ac > 0$ , then

$$y = c_1 e^{r_1 x} + c_2 e^{r_2 x}.$$

2. If  $r_1 = r_2$ , because  $b^2 - 4ac = 0$ , then

$$y = c_1 e^{r_1 x} + c_2 x e^{r_1 x}.$$

3. If  $r_1 = \alpha + i\beta$  and  $r_2 = \alpha - i\beta$ , because  $b^2 - 4ac < 0$ , then

$$y = e^{\alpha x} (c_1 \sin(\beta x) + c_2 \cos(\beta x)).$$

## 2.10.2 Nonhomogeneous ODEs with Constant Coefficients

A nonhomogeneous differential equation with constant coefficients has the form:

$$a \frac{d^2 y}{dx^2} + b \frac{dy}{dx} + cy = G(x), \quad (2.10.3)$$

where  $a$ ,  $b$  and  $c$  are constant coefficients and  $G(x)$  is a continuous function.

To solve equations with this form there are two primary methods. The first one is called *The Method of Undetermined Coefficients* and the second is called *The Method of Variation of Parameters*. Because the first method is conceptually quite simple (although it does have a few oddities that you need to look out for) and can be extended to include matrix solution methods, discussed in Section ??, it will be the method we use here.<sup>6</sup> Also worth noting is that the explanations given here are geared toward solving specific problems and for a thoroughly general exposition on these topics the reader is encouraged to turn to [?] and [?].

<sup>6</sup>Variation of Parameters also extends nicely to include matrix solution methods.

Regardless of which method you use, it is important to know that a general solution to a nonhomogeneous differential equation has the form:

$$y(x) = y_c(x) + y_p(x),$$

where  $y_c(x)$  and  $y_p(x)$  are linearly independent  $y_c(x)$  is the general solution to the related homogeneous equation

$$a \frac{d^2 y}{dx^2} + b \frac{dy}{dx} + cy = 0,$$

and is called the *complementary equation*.  $y_c(x)$  can be derived using the method given in Section ?.  $y_p(x)$  is the particular solution to Equation ?? and can be derived using *Undetermined Coefficients* or *Variation of Parameters*.

The method of undetermined coefficients simply sets  $y_p(x)$  to a general polynomial of the same degree as  $G(x)$  and then this is substituted into the differential equation to solve for the coefficients. That is to say, if  $G(x) = x^2$ , then we start with  $y_p(x) = Ax^2 + Bx + C$ . If  $G(x) = e^{3x}$ , then  $y_p(x) = Ae^{3x}$ , and if  $G(x) = \sin x$ , then  $y_p(x) = A \cos x + B \sin x$ . The only significant catch is when  $G(x)$  has the same form as  $y_c(x)$ . The effect of this would make  $y_c(x)$  and  $y_p(x)$  linearly dependent. In this case, we simply multiply  $y_p(x)$  by  $x$ , or  $x^2$ , whichever makes  $y_p(x)$  linearly independent from  $y_c(x)$ .

### Example 2.10.2.1

To find a solution to the equation:

$$\frac{d^2 y}{dx^2} + \frac{dy}{dx} - 2y = \sin x, \quad (2.10.4)$$

we first find a solution to the complementary equation

$$\frac{d^2 y}{dx^2} + \frac{dy}{dx} - 2y = 0.$$

We do this by writing out the characteristic equation

$$r^2 + r - 2 = 0$$

which has roots  $r_1 = 1$  and  $r_2 = -2$ . By using the rules found in Section ?? we know that the general solution to the complementary equation is

$$y_c(x) = c_1 e^x + c_2 e^{-2x}.$$

To determine  $y_p(x)$ , we start with

$$y_p(x) = A \cos x + B \sin x.$$

Differentiating gives us:

$$\begin{aligned}y_p'(x) &= -A \sin x + B \cos x \\y_p''(x) &= -A \cos x - B \sin x.\end{aligned}$$

We now substitute these equations into Equation ?? and attempt to solve for the coefficients. Thus,

$$\begin{aligned}(-A \cos x - B \sin x) + (-A \sin x + B \cos x) \\- 2(A \cos x + B \sin x) = \sin x\end{aligned}$$

which reduces to

$$(-3A + B) \cos x + (-A - 3B) \sin x = \sin x.$$

We can now solve for  $A$  and  $B$  with the following equations:

$$\begin{aligned}-3A + B &= 0 \\-A - 3B &= 1.\end{aligned}$$

Thus,  $A = -\frac{1}{10}$  and  $B = -\frac{3}{10}$  and

$$y_p(x) = -\frac{1}{10} \cos x - \frac{3}{10} \sin x.$$

Combining  $y_c(x)$  and  $y_p(x)$  gives us a general solution to Equation ??:

$$y(x) = c_1 e^x + c_2 e^{-2x} - \frac{1}{10} \cos x - \frac{3}{10} \sin x.$$

||

### Example 2.10.2.2

To find a solution to the equation:

$$\frac{d^2 y}{dx^2} + y = \sin x \quad (2.10.5)$$

we start by finding a general solution to the complementary equation

$$\frac{d^2 y}{dx^2} + y = 0,$$

which has the characteristic equation

$$r^2 + 1 = 0$$

and roots  $r_1 = i$  and  $r_2 = -i$ . Thus,

$$y_c(x) = c_1 \cos x + c_2 \sin x$$

To determine  $y_p(x)$ , we first use the equation

$$y_p(x) = A \cos x + B \sin x,$$

but this has the same form as  $y_c(x)$ , and thus, we must multiply it by  $x$ . Thus,

$$y_p(x) = Ax \cos x + Bx \sin x,$$

and

$$\begin{aligned}y_p'(x) &= A \cos x - Ax \sin x + B \sin x + Bx \cos x \\y_p''(x) &= -A \sin x - A \sin x - Ax \cos x \\&\quad + B \cos x + B \cos x - Bx \sin x \\&= -2A \sin x - Ax \cos x + 2B \cos x - Bx \sin x.\end{aligned}$$

Substitution into Equation ?? gives us

$$-2A \sin x + 2B \cos x = \sin x$$

which results in the following system of equations:

$$-2A + 0B = 10A + 2B = 0,$$

thus,  $A = -\frac{1}{2}$  and  $B = 0$ , and

$$y_p(x) = -\frac{1}{2} x \cos x.$$

Combining  $y_c(x)$  and  $y_p(x)$  give us the following general solution to Equation ??:

$$y(x) = c_1 \cos x + c_2 \sin x - \frac{1}{2} x \cos x.$$

||

## 2.10.3 Using Matrix Algebra

Differential equations can be solved using matrix algebra. The advantages to this method are that second, third, fourth and even higher order differential equations can be all solved as first order differential equations and that when forced to find a numerical solution (as opposed to the analytic solutions), most computer tools expect the input to be in matrix format.

## 2.11 Numerical Approximations to ODEs

Finding an analytical solution to a differential equation is not always a practical option. Numerical approximations lead to solutions that are much more readily available, however, there are a number of issues related to the approximation that should be understood. The trick to constructing a viable numerical

solution of a differential is identifying a reliable approximation of the derivative and then selecting a step size that results in both a stable and physically meaningful solution.

**THE MAIN IDEA - Crafting approximations:**

Numerically solving a differential equation requires an initial condition (the point where the solution starts) and an algorithm for extending the solution. Extending the solution can be visualized as an exercise with the Taylor series. The idea is to expand the solution space around the initial condition, and use the Taylor series to guide the approximation of the solution.

Errors acquired during the construction of a numerical solution arise from two sources: roundoff and truncation. Roundoff error arises from the limited precision of computer arithmetic. The problem is compounded in that the binary representation of many fractions is irrational, enhancing the effects of the roundoff error. For example, 1/10 is irrational in the binary system whereas 1/8 is rational. For this reason, we often choose discretization intervals that are powers of 1/2 instead of powers of 1/10.

The second source of error is called truncation error. This error arises when we make discrete approximations of continuous functions. This error can be, to a certain extent, limited by making the step-sizes in the discrete function as small as possible. The Taylor series, which provided a means for creating approximate functions, also allows us to evaluate the truncation error. We often evaluate the quality of a numerical solution by estimating the error incurred with our functional approximations.

**Example 2.11.0.1 (Forward Euler Method)**

We start with a simple first order initial value problem (IVP)

$$\frac{dy}{dt} = f(y, t), \text{ and } y(t_0) = y_0 \quad (2.11.1)$$

where  $f(y,t)$  is some linear or non-linear function of  $y$  and  $t$  and  $y(t_0)$  is the initial value of  $y$  at time  $t_0$ .

Finding approximate values for points other than  $t_0$  is simply a matter of figuring out what step size,  $h$ , of the independent variable,  $t$  to use. Figure ?? shows the idea of extending the solution using a forward Euler solution.

The strategy is to start at a point and extend the solution with a step size,  $h$ . So we start at  $(y_0, t_0)$ ,




Figure 2.11.1: The general idea for extending a solution using a forward Euler method.

compute the amount of  $y$  we must add,  $f(y_0, t_0) * h$  and then use this point to repeat the algorithm:

1.  $t_{i+1} = t_i + h$
2.  $y_{i+1} = y_i + h * f(y_i, t_i)$

This is called the *Forward Euler* method because the right hand side,  $f(y, t)$  is evaluated at the initial point. Now, is this a good method? We'll now look at a Taylor expansion of our method, and estimate the error and the stability of the solution. By stability we mean will the numeric solution reflect the analytic solution for arbitrary conditions (like step size and nature of the function,  $f(y, t)$ ).

Now, using the Taylor series, Equation ??, we will expand  $y(t)$  around  $y(t_i)$  and assess the truncation error:

$$y(t_{i+1}) = y(t_i) + hy'(t_i) + \frac{h^2}{2}y''(t_i) + O(h^3). \quad (2.11.2)$$

The three left terms represent the truncated Taylor series (Euler's Method) while the right two terms represent the local truncation error. We are assuming that  $h < 1$  and thus, the  $O(h^3)$  term indicates that all terms that follow in the series will be less than  $h^3$ . We see that the method is accurate to order  $h$ , and not the higher order terms ( $h^2$  or  $h^3$ ), so we say that its first order accurate.

In order to determine how stable this method is for approximating an ordinary differential equation, we simply try to determine whether or not it will converge or diverge for large values of  $i$ . For example, we will use a simple function:  $f(y, t) = \lambda y$  so that

$$f(y, t) = \frac{dy}{dt} = \lambda y, \text{ and } y(0) = y_0.$$

Using this to replace the function,  $f(y, t)$ , in Equation ??, we have

$$y_{i+1} = y_i + h\lambda y_i = (1 + h\lambda)y_i.$$



Now, if you start at the initial value,  $y_0$ , and repeatedly apply the above equation, you'll find that

$$\begin{aligned} y_1 &= (1 + h\lambda)y_0 \\ y_2 &= (1 + h\lambda)(1 + h\lambda)y_0 \\ &\vdots \\ y_{i+1} &= (1 + h\lambda)^{i+1}y_0. \end{aligned}$$

We call the multiplier,  $1 + h\lambda$  the amplification factor. Clearly if  $h\lambda > 0$  then the solution will blow up after a few steps. Thus, the solution is only stable when  $\lambda < 0$ , since the step size,  $h$ , can never be negative itself. ||

### Example 2.11.0.2 (Backward Euler Method)

The Forward Euler method approximates the points  $y_{i+1}$  by starting from some initial point,  $y_0$  and moving to the right using the derivative as calculated at  $y_i$ . An alternative to this is to start from  $y_0$  and move to the right using the derivative as calculated at  $y_{i+1}$ . That is, we evaluate the function,  $f(y, t)$  at  $y_{i+1}, t_{i+1}$ . To do this, we rewrite the second formula in the numerical approximation as

$$y_{i+1} = y_i + hf(y_{i+1}, t_{i+1}). \quad (2.11.3)$$

Now we let  $f(y, t) = \lambda y$  and solve for  $y_{i+1}$  and find that

$$y_{i+1} = y_i + h\lambda y_{i+1},$$

and thus,

$$y_{i+1} = \left( \frac{1}{1 - h\lambda} \right) y_i.$$

Starting at the initial condition,  $y_0, t_0$  we see that

$$\begin{aligned} y_1 &= \left( \frac{1}{1 - h\lambda} \right) y_0 \\ y_2 &= \left( \frac{1}{1 - h\lambda} \right) \left( \frac{1}{1 - h\lambda} \right) y_0 \\ &\vdots \\ y_{i+1} &= \left( \frac{1}{1 - h\lambda} \right)^i y_0, \end{aligned}$$

and the amplification factor is now  $< 1$  for all values of  $\lambda < 0$  (again, since the step size,  $h$ , can never be negative). Thus, this method only provides a stable solution when  $\lambda < 0$ . This method is referred to as an implicit method, since the function is evaluated at a solution point, yet to be determined. For all

linear functions,  $f$ , a nice iterative equation can be determined. When  $f$  is non-linear, then typically the function is linearized (with a Taylor series, of course) and each step in the solution is iterated until the function is well approximated by the Taylor approximation. ||

### Example 2.11.0.3 (Leap Frog Method)

The error associated with the simple Euler method can be improved by realizing that for both the forward (explicit) and backward (implicit) Euler methods, there is an asymmetry between approximating the derivative and evaluating the function. For the explicit method, the function (right hand side) is evaluated at the left side of the derivative while for the implicit method,  $f(y, t)$  is evaluated at the right side of the derivative. The price for this approximation is that we now need to know two values of  $y$   $y_i$  and  $y_{i-1}$  before we can extend the solution to the next point. We can also evaluate  $f(y, t)$  at the midpoint of the derivative, and this is called the *Leap Frog Method*.

$$\frac{dy}{dt} = \frac{y_{i+1} - y_{i-1}}{2h}$$

so that

$$y_{i+1} - y_{i-1} = 2hf(y_i, t_i)$$

As before, we can evaluate the accuracy of this strategy by comparing to the Taylor expansion of the function. We first make a Taylor approximation for both  $y_{i+1}$  and  $y_{i-1}$ :

$$y(t_{i+1}) = y(t_i) + hy'(t_i) + \frac{h^2}{2}y''(t_i) + O(h^3),$$

and

$$y(t_{i-1}) = y(t_i) - hy'(t_i) + \frac{h^2}{2}y''(t_i) - O(h^3).$$

Now, subtracting the two series, we see that the the 2nd order terms cancel. Thus, this method is 2nd order accurate. ||

### Example 2.11.0.4 (Runge-Kutta Method)

The solution accuracy can be further improved by building a strategy around the leap-frog method and the forward Euler method. This mixture results in a Runge-Kutta method that has even more accuracy. The half step is computed with the forward Euler method and then the full step is computed with the leap-frog method.

$$y_{i+1/2} = y_i + \frac{h}{2} f(y_i, t_i) \quad (2.11.4) \quad > \text{more fhn\_trigger.m}$$

$$y_{i+1} = y_i + h f(y_{i+1/2}, t_{i+1/2}) + O(h^3) \quad (2.11.5) \quad \text{t} = \text{linspace}(0, 50, 400);$$

Like the leap frog method, this method is 2nd order accurate. The Runge-Kutta method applies to a class of methods where intermediate steps are taken. For example, we can make 4 evaluations and make a procedure that is 4th order accurate:

$$k_1 = f(y_i, t_i) \quad (2.11.6) \quad \text{plot}(t, y);$$

$$k_2 = f\left(y_i + \frac{k_1}{2}, t_i + \frac{h}{2}\right) \quad (2.11.7) \quad \text{z} = [\text{t}' \text{ y}];$$

$$k_3 = f\left(y_i + \frac{k_2}{2}, t_i + \frac{h}{2}\right) \quad (2.11.8) \quad \text{save -ascii plot.dat z};$$

$$k_4 = f(y_i + k_3, t_i + h) \quad (2.11.9)$$

$$y_{i+1} = f(y_i, t_i) + \frac{k_1}{6} + \frac{k_2}{3} + \frac{k_3}{3} + \frac{k_4}{6} + O(h^5) \quad (2.11.10)$$

This procedure is a very popular procedure, and will usually do you right. It is safe, accurate, and except for really wierd models, will provide reliable solutions.

||

Fortunately, octave (matlab) has some very nice tools for solving odes for us. Here are two little scripts just to indicate what is needed. The first segment defines the ode as a function, in this case  $du/dt = u(1-u)(u-a) + \text{stim}$  where  $\text{stim}$  is the stimulation that forces the system to switch from one stable state to another:

```
> more trigger.m

function xdot=trigger(x,t)

% Trigger threshold for stim = 0.25 is 1.039 time units

xdot = zeros(1,1);
a = 0.25;
stim = 0;

if t < 1.039
    stim = 0.25;
end

xdot(1) = x(1)*(1.0-x(1))*(x(1)-a) + stim;

endfunction
```

To to actually solve the ode, we execute the following piece of code within matlab which calls the ode solver: `lsode` and then plots the results

## 2.12 Markov and Stochastic Models

# Chapter 3

## How to ask questions about a model

### 3.1 Numerical Experiments

One of the easiest ways to explore a model is to simply make numerical experiments. These experiments can be as simple as plotting the model response as a function of the independent variable(s). Over the years we have accumulated a toolbox filled with a number of useful command-line tools for performing these experiments. The toolbox is available as a tar file (analyze.tgz) for use on any system with a C compiler.

A simple example of these tools in action is as follows:

```
shell> series 0 10 .1 | \
> tf x1 2.5*exp(-3.5*x1) | xmgrace -pipe
```

The first part of the pipe generates the independent variable values, the second component transforms the independent variable into the model function, in this case, an exponential. The third component plots the data using grace (or the graphing program of your choice).

If the model is not an algebraic model, but rather described by the solution to an ordinary or partial differential equation, then often a simple Euler integration and its graph is adequate to give a general idea of what the model is doing. The Section ?? gives an overview of this and other numerical methods.

The bottom line is always to plot the results, never simply look at a list of numbers. It is surprising what a graphical representation will reveal that numbers hide.




Figure 3.2.1: A histogram of exam scores

### 3.2 Formal Questions of a model and data: Statistical Hypothesis Testing

#### 3.2.1 Probability Distributions

Imagine a class with 100 students and at the end of the semester they all took the final exam. After the teacher graded the exam he noticed that a lot of students scored between 70 and 75 percent. A slightly smaller group of students must have studied a little harder as they got between 75 and 90 percent, and a similar sized group probably spent more time at the local pub because they scored between 55 and 70 percent. Even rarer were the groups that aced or totally bombed the test.

If the teacher wanted to visualize how the grades on his final exam were distributed, he could draw out a histogram (see Figure ??) that would show the number of exams with scores that fell into different ranges. If he then scaled each column in the his-

togram by dividing by the total number of exams (in this case, 100), the histogram would also give the teacher a rough estimate of the probability of picking an exam at random with a score between 70 and 90. If the teacher wanted to know what the probability the exam score would be between 70 and 90, he could simply add the columns that represented that range together to give himself a general idea.

If the teacher fit a curve to the scaled version of the histogram, so that the total area under the curve was equal to 1, he could use integration to estimate the probability of an exam having a score between two points. The smaller the area between two points, the lower the probability that an exam will have that score. The larger the area between two points, the greater the probability of randomly selecting an exam in that range. These probabilities are confirmed by the original histogram.

At this point, you may be wondering why the teacher would want to use a fitted curve instead of his original histogram to determine probabilities. The reason for this is that it is easier to compare curves, and thus, use them to answer questions. If the teacher fit curves to several years' worth of histograms, he could use them to determine if having the study session two days before the exam helped or not.

### 3.2.2 Comparing Two Samples: Clarifying Variance

Imagine that the teacher mentioned in the previous section was indeed attempting to determine if the study session he held two days before the exam helped or not. One potential pitfall in simply visually comparing two graphs would be that an improvement in scores could be due to several factors, not just to the study session. For example, this year's class, as a whole, could have been luckier in the multiple choice section than the previous year's. Is there any way to characterize how lucky a class would need to be in order to perform much better? If there was (and there is), then perhaps the teacher could interpret the fact that the class would have to be amazingly lucky to perform as well as it did as not just luck, but due to the study session. In a sense, the fundamental idea of statistics is to try to determine if change is a result of chance or due to a specific reason. In order to do this, we must consider variability (see Figures ??, ?? and ??).

The variance of a model, notated with  $\sigma^2$ , quantifies how spread out the data is. In Figures ??, ??




Figure 3.2.2: Two normal pdfs. One with mean  $\mu = 0$ , and the other with  $\mu = 4$ . Both have the same variance,  $\sigma^2 = 1$ . Notice that with a small amount of variance, the two graphs hardly overlap and it is easy to distinguish how one is fundamentally different from the other.




Figure 3.2.3: Two normal pdfs. One with  $\mu = 0$  and the other with  $\mu = 4$ . Both have the same variance,  $\sigma^2 = 2$ . Notice how with the increased variance, the two curves overlap each other more than in Figure ??.




Figure 3.2.4: Two normal pdfs. One with  $\mu = 0$  and the other with  $\mu = 4$ . Both have the same variance,  $\sigma^2 = 4$ . Compare the variance and the overlap with Figures ?? and ?. The more overlap there is, the harder it is to determine if they are fundamentally different.

and ?? we can see how larger values for  $\sigma^2$  result in shorter peaks and a broader distribution of the data. If we can determine the variance for a model (usually it is estimated from the data), then we can use this information to help calculate the significance of the difference between two samples. Typically, this is done by calculating difference between the two sample means and dividing by the square root of the variation. That is,

$$\frac{\hat{\mu}_A - \hat{\mu}_B}{\sqrt{s^2}}, \quad (3.2.1)$$

where  $\hat{\mu}_A$  is the mean of the  $A$  data set,  $\hat{\mu}_B$  is the mean of the  $B$  data set and  $s^2$  is an estimate of the variation in the distributions that the data came from,  $\sigma^2$ . We will talk more about how to calculate  $s^2$  shortly.

If you think about it briefly, you will see that the smaller the variation, the more significant the difference will be. Another way to think about it is that, for data sets with large amounts of variation, the difference between the two means must be greater in order to avoid being in the area of overlap between the two distributions (Figure ??).

#### Example 3.2.2.1 (The Study Session)

For example, if the average exam score for the year without the study session was 75%, thus,  $\hat{\mu}_A = 75$  and the and for the year with the study session  $\hat{\mu}_B = 79$ , then the difference the two years is,  $\hat{\mu}_A - \hat{\mu}_B = 4$ . If the estimated variation in the distributions that the data came from is  $s^2 = 1$ , then we will have distributions like that seen in Figure ?? and the difference

between the two means would be quite clear. However, if  $\sigma^2 = 4$ , then the difference would be scaled by Equation ?? to be only 2, and not as significant.

||

Calculating  $s^2$ , an estimate of  $\sigma^2$  is quite simple. We simply average the squared differences between each observation and the mean. That is,

$$s^2 = \frac{\sum(x_i - \bar{x})^2}{n - 1}. \quad (3.2.2)$$

The reason we square each difference is that we do not want positive deviations from the mean negating negative deviations.

To summarize the process of statistical analysis, here is a list of general steps:

1. Take a bunch of measurements.
2. Make a histogram of the measurements. From this we can take a guess at the type of distribution that the data came from. In this case, the histogram looked fairly symmetrical with a single hump in the middle and this shape is often modeled with a *normal distribution*. Other shapes are better modeled with other distributions (see Figure ??).
3. Estimate means and variances from the data and use them to compare different distributions.

### 3.3 What Statistical Power Means

*Power* is a term that is used quite frequently to describe statistical tests. As is often the case, the word has a rather specific definition which we will attempt to describe here. Due to their close relation to the definition of power, we will also briefly describe the various types of errors that statistical tests can make. Thus,

$\alpha$  = the probability you will reject  $H_0$  when it is true. This type of error is called Type I Error.

$\beta$  = the probability you will accept  $H_0$  when it false. This type of error is called Type II Error.

Power =  $1 - \beta$ , the probability the test will reject  $H_0$  when it is false. Thus, the more power, the higher probability of correctly rejecting  $H_0$ .




Figure 3.3.1: The predicted distribution of  $\bar{X}$  given by the null hypothesis,  $H_0 : \mu \leq 0$ , is depicted in the top graph. The bottom graph shows the true distribution of  $\bar{X}$ , since  $\mu = 1$ . With the current sample size, the variation in  $\bar{X}$ , our estimator for  $\mu$ , is great enough to make it more than likely that we will fail to reject  $H_0$  even though it is false.

You can increase power by increasing the sample size,  $n$ , for the test. This is because the larger sample size will decrease the variance of the estimated parameters. For example, consider  $\bar{X}$  as an estimate of  $\mu$ . By the central limit theorem, the variance of  $\bar{X}$ , where  $EX = \mu$  and  $\text{Var}(X) = \sigma^2$  for independent and identically distributed samples from *any* distribution, is approximately  $\sigma^2/n$ , which gets smaller as  $n$  gets larger.

An example of this is shown in Figures ?? and ??.

### 3.3.1 Numerical Approximations of Power

There are two ways to determine the power of a statistical test, analytical and numerical (that is, experimentally) and quite often, the numerical method is a more practical approach. To determine the power of a test given a sample size  $n$ , you need only to follow these steps:

1. Generate  $n$  simulated data points (see Section ?? for details on how this can be accomplished).
2. Apply your test to the data.
3. Determine if your test correctly rejected the null hypothesis.
4. Return to Step ?? and repeat many, many times.
5. The power of your test is:

$$\text{Power} = \frac{\# \text{ of times } H_0 \text{ was correctly rejected}}{\text{Total } \# \text{ of tests}}$$




Figure 3.3.2: The predicted distribution of  $\bar{X}$  given by the null hypothesis,  $H_0 : \mu \leq 0$ , is shown in the top graph. The bottom graph depicts the true distribution of  $\bar{X}$ , since  $\mu = 1$ . However, compared with Figure ??, the sample size has been increased enough to reduce the variation in the parameter estimate by one half. This makes it more likely that our test will reject  $H_0$ , and thus, the test has more power.

## 3.4 What a 95% Confidence Interval Is

The concept of a 95% Confidence Interval (95% CI) is one that is somewhat elusive. This is primarily due to the fact that many students of statistics are simply required to memorize its definition without fully understanding its implications. Here we will try to cover both the definition as well as what the definition actually implies.

The definition that students are required to memorize is:

If the procedure for computing a 95% confidence interval is used over and over, 95% of the time the interval will contain the true parameter value.

Students are then told that this definition does not mean that an interval has a 95% chance of containing the true parameter value. The reason that this is true, is because a 95% confidence interval will either contain the true parameter value of interest or it will not (thus, the probability of containing the true value is either 1 or 0). However, you have a 95% chance of creating one that does. In other words, this is similar to saying, "you have a 50% of getting a heads in a coin toss, however, once you toss the coin, you either have a head or a tail". Thus, you have a 95% chance of creating a 95% CI for a parameter that contains

the true value. However, once you've done it, your CI either covers the parameter or it doesn't.

### 3.5 P-Values

In publications, you will often times see p-value reported as the result of some statistical test. A p-value is the probability of an event (or series of events) coming up that would create a statistic with an extreme value<sup>1</sup> than the one you derived, assuming your model under the null hypothesis is correct.

Regardless of the type of model you are assuming describes the source of the data under the hypothesis, you can create what are called one-sided tests. With these tests, there are two typical hypotheses that people make about the mean of the underlying model. One type of hypothesis is that the mean is less than some value. For example, you might propose that the mean is less than zero, or  $H_0 : \mu < 0$ . Alternatively, the hypothesis might be that the mean is greater than zero, or  $H_0 : \mu \geq 0$ .

For the first type of one-sided hypothesis, the p-value is defined as:

$$\begin{aligned} \text{p-value} &= \Pr(x \geq \text{your statistic}) \\ &= \int_{\text{your statistic}}^{\infty} f(x)dx, \end{aligned}$$

where  $f(x)$  is the probability distribution you are assuming the data came from, and *your statistic* is the value derived from a function of the data (for example, the mean of the data). This is illustrated in Figure ???. Since our hypothesis is  $H_0 : \mu \leq 0$ , the smaller the mean of the data is (and thus, the smaller the p-value), the more likely we will reject the proposed model.

The second type of one-sided hypothesis, where we are testing to see if the mean is greater than some value, is very similar. The only difference is that we integrate in the other direction. That is:

$$\begin{aligned} \text{p-value} &= \Pr(x \leq \text{your statistic}) \\ &= \int_{-\infty}^{\text{your statistic}} f(x)dx. \end{aligned}$$

This is illustrated in Figure ???.

If the type of model you are assuming describes the source of the data is symmetric (like the distributions in Figures ??? and ???) you can create what

<sup>1</sup>That is, a greater value, or lesser value, or both, depending on the model and the type of hypothesis you are testing. The details of this will be explained in the next few paragraphs.




Figure 3.5.1: The p-value for a one-sided statistic where we are testing  $H_0 : \mu \leq 0$ .




Figure 3.5.2: The p-value for a one-sided statistic where we are testing  $H_0 : \mu \geq 0$ .




Figure 3.5.3: The p-value for a two-sided statistic where we are testing  $H_0 : \mu = 0$ .

are called two-sided tests. In this case your typical null hypothesis is that the mean is equal to a certain value. For example, you might propose the hypothesis  $H_0 : \mu = 0$ . Thus, if the mean of your data is much larger or much smaller than zero, then you have good reason to reject  $H_0$ . In this case, the p-value is defined as:

$$\begin{aligned} \text{p-value} &= \Pr(x \geq |\text{your statistic}|) \\ &\quad + \Pr(x \leq -|\text{your statistic}|) \\ &= 2 \times \Pr(x \geq |\text{your statistic}|) \\ &= 2 \int_{|\text{your statistic}|}^{\infty} f(x) dx. \end{aligned}$$

This is illustrated in Figure ??.

Obviously, the smaller the p-value, the less likely an event as rare or rarer will take place. Often times the model proposed by the null hypothesis,  $H_0$ , is rejected if the p-value is less than 0.05. That is to say, it is assumed that the proposed model does not explain the data if the p-value is less than 0.05.

### 3.6 Degrees of Freedom

Estimates of parameters can be based upon different amounts of information. The number of independent pieces of information that go into the estimate of a parameter is called the degrees of freedom (df). In general, the degrees of freedom of an estimate is equal to the number of independent values that go into the estimate minus the number of parameters estimated

as intermediate steps in the estimation of the parameter itself.

For example, for a random sample of  $n$  independent data points, if the sample mean,  $\bar{X}$  is estimated using the standard formula  $1/n \sum x_i$ , then the degrees of freedom for  $\bar{X}$  is  $n$ . This is because  $\bar{X}$  uses all of the independent values from the sample and does not rely on any other parameter estimates in its calculation. However, if the variance,  $s^2$ , is estimated using the standard formula  $1/(N - 1) \sum (x_i - \bar{X})^2$ , then the degrees of freedom is equal to the number of independent values ( $n$ ) minus the number of parameters estimated as intermediate steps (one,  $\bar{X}$ ) and is therefore equal to  $N-1$ .

### 3.7 Chi-Square Goodness of Fit Test

In the study of genetics one frequently runs into situations that are resolved using what is called a Chi-Square Goodness of Fit Test. This is a test that is particularly adept at determining how well a model fits observed data. It allows us to evaluate how “close” the observed values are to those which would be expected given the model in question. Here is a brief explanation of how and why the Chi-Square Goodness of Fit Test is effective in these situations.<sup>2</sup>

In general, the chi-square test statistic has the form:

$$\chi^2 = \sum \frac{(\text{observed} - \text{expected})^2}{\text{expected}}, \quad (3.7.1)$$

and if  $\chi^2$  is large, than the model is a poor fit to the data. Before we get into the details of the theory behind this statistic, let’s begin with a short example of how it is used.

#### Example 3.7.0.1 (A Fair Coin?)

Imagine trying to determine if a coin is fair or not. If the coin is fair, than the probability of getting heads is  $p_1 = 0.5$  and the probability of getting tails is  $p_2 = 0.5$ , other wise  $p_1 \neq 0.5$  and  $p_2 \neq 0.5$ . It is important to note that since the coin has only two sides,  $p_2 = 1 - p_1$ . While this equality may seem obvious, it will be useful when we are determining the degrees of freedom for our test. If we tossed the coin 100 times,

<sup>2</sup>A lot of the material in this section was plagiarized from the web page: <http://www.stat.yale.edu/Courses/1997-98/101/chigf.htm>, author unknown.



we would expect to get heads  $100 \times 0.5 = 50$  times. We know, however, that even though the probability of getting heads is 0.5, there is a chance that we might get a few more or a few less than 50 heads in 100 tosses. The question is, how much variation in the number of heads will we allow before we are confident in rejecting the hypothesis that  $p = 0.5$ . This is where the Chi-Square Goodness of Fit test comes in handy.

In order to test the hypothesis that the coin is fair, you toss the coin 100 times and observe that it landed on heads 38 times. From this data alone, we are able to determine that the coin must have landed on tails 62 times and we note this in Table ??.

	Observed	Expected
Heads	38	50
Tails	62	50

Table 3.7.1: Both observed and expected results of 100 coin tosses.

With this data in our hands, we can compute a  $\chi^2$  test statistic and use it to determine the fairness of the coin. That is,

$$\begin{aligned} \chi^2 &= \frac{(38 - 50)^2}{50} + \frac{(62 - 50)^2}{50} \\ &= \frac{144}{50} + \frac{144}{50} \\ &= 5.76. \end{aligned}$$

We can now see where this value lies in a  $\chi^2$  distribution. If it is in the tail of the distribution, then the probability of getting 37 heads using a fair coin would appear to be a very rare event. If it is in the middle of the distribution, then it might be quite common to obtain 38 heads in 100 tosses from a fair coin.

In order to examine our value in the context of a  $\chi^2$  distribution we must specify which one by determining its degrees of freedom. We calculate the total degrees of freedom by looking at the total number of parameters in our model, 2 ( $p_1$  and  $p_2$ ), and subtracting 1 because  $p_2$  is not independent from  $p_1$  since  $p_2 = 1 - p_1$ . Thus, we must see how much area is under the curve of a  $\chi^2_1$  distribution (the subscript 1 indicates the degrees of freedom) from 5.75 to  $\infty$ . We can do this easily using Octave:

```
octave:1> 1 - chisquare_cdf(5.76, 1)
ans = 0.016395
```




Figure 3.7.1: The area under the  $\chi^2_1$  graph that represents the p-value, the probability our hypothesis that the coin is fair is correct. Since the p-value/area is so small (1.6 percent) we will reject our hypothesis.

The probability that a value of 5.76 or larger would come from the  $\chi^2_1$  distribution is less 0.016395, which is very small (see Figure ??). Much smaller than the standard 5 percent used as a cutoff to determine whether we should accept 5.76 as coming from the  $\chi^2_1$  distribution. Thus, we will reject the hypothesis that this coin is fair. ||

### 3.7.1 Why It Works

Consider a binomial random variable  $Y \sim \text{Bin}(n, p)$  with mean  $\mu_Y = np$  and variance  $\sigma_Y^2 = np(1 - p)$ . From the Central Limit Theorem, we know that  $Z = (Y - \mu_Y)/\sigma_Y$  has an approximately a standard Normal(0,1) distribution for large values of  $n$ . Since the square of a standard normal random variable has a chi-square distribution with one degree of freedom,  $Z^2$  is approximately  $\chi^2_1$ .

Now consider the random variable  $Y_1$  which has a binomial( $n, p_1$ ) distribution and let  $Y_2 = n - Y_1$  and  $p_2 = 1 - p_1$ . Then

$$\begin{aligned} Z^2 &= \frac{(Y_1 - np_1)^2}{np_1(1 - p_1)} \\ &= \frac{(Y_1 - np_1)^2(1 - p_1) + (Y_1 - np_1)^2(p_1)}{np_1(1 - p_1)} \\ &= \frac{(Y_1 - np_1)^2}{np_1} + \frac{(Y_1 - np_1)^2}{n(1 - p_1)}, \end{aligned}$$

and since

$$(Y_1 - np_1)^2 = (n - Y_2 - n + np_2)^2 = (Y_2 - np_2)^2,$$

we have

$$Z^2 = \frac{(Y_1 - np_1)^2}{np_1} + \frac{(Y_2 - np_2)^2}{np_2},$$

where  $Z^2$  has a chi-square distribution with 1 degree of freedom.

In general, for  $k$  random variables  $Y_i$ , where  $i = 1, 2, \dots, k$ , with corresponding expected values  $np_i$ , a statistic measuring the “closeness” of the observations to their expectations is the sum:

$$\frac{(Y_1 - np_1)^2}{np_1} + \frac{(Y_2 - np_2)^2}{np_2} + \dots + \frac{(Y_k - np_k)^2}{np_k},$$

which has a chi-square distribution with  $k - 1$  degrees of freedom. This is because we know that the sum of all of the probabilities,  $p_1, \dots, p_k$ , must equal 1, and thus we can derive  $p_k$  by subtracting the first  $k - 1$  probabilities from 1.

### Example 3.7.1.1 (Allele Frequencies)

The population is said to be in Hardy-Weinberg equilibrium for a given gene if it is:

1. Stable with respect to the allele and genotype frequencies of interest. That is, allele frequencies do not change from generation to generation.
2. The genotype frequencies in the progeny produced by random mating among parents is determined solely by the allele frequencies of the parents.

In other words, if, for a particular gene **A** with alleles  $A_1$  and  $A_2$ , and the allele frequencies in the parents are  $f(A_1) = p$  and  $f(A_2) = q$  (and thus  $p + q = 1$  or  $q = 1 - p$ ), then the percentage of offspring with the genotype  $A_1A_1 = p^2$ ,  $A_1A_2 = 2pq$  and  $A_2A_2 = q^2$ .

Genotype	Observed
$A_1A_1$	22
$A_1A_2$	216
$A_2A_2$	492

Table 3.7.2: Observed genotypes at the MN blood group gene locus for individuals in a human population. *Source:* Plagiarized from Michael D. Purugganan, class notes.

Given the data in Table ??, we can calculate the observed allele frequencies. That is,

$$p = \frac{(22 + 216/2)}{730} = 0.178,$$

and

$$q = 1 - p = 0.822.$$

With values for  $p$  and  $q$ , we can now calculate how many individuals with each class of genotype we would expect if the population was in Hardy-Weinberg Equilibrium. The results of this calculation are in Table ??.

Genotype	Observed	Expected
$A_1A_1$	22	23.14
$A_1A_2$	216	213.60
$A_2A_2$	492	493.26

Table 3.7.3: Both observed and expected genotypes at the MN blood group gene locus for individuals in a human population.

Now that we have both observed and expected values for each class of genotype, we can calculate a chi-square test statistic. That is,

$$\begin{aligned} \chi^2 &= \frac{(22 - 23.14)^2}{23.14} + \frac{(216 - 213.60)^2}{213.60} \\ &\quad + \frac{(492 - 493.26)^2}{493.26} \\ &= 0.086 \end{aligned}$$

Now all we need to do is compare this value to that from a chi-square distribution. The trick, however, is determining how many degrees of freedom there are. Here we have three different categories, or genotypes, and each one has an associated probability of membership. However, two of these probabilities are dependent on one of them. That is, since  $q = 1 - p$  the probability of having the genotype  $A_1A_2 = 2pq = 2p(1 - p)$  and the probability of having the genotype  $A_2A_2 = q^2 = (1 - p)^2$ . Thus, since there is only one linearly independent probability, the degree of freedom is 1.

We can now use Octave to determine the probability our hypothesis is correct:

```
octave:2> 1 - chisquare_cdf(0.086, 1)
ans = 0.76933
```

So, since we usually fail to reject the hypothesis that the data comes from our model if the probability is more than 5 percent (and in this case it is 77 percent, see Figure ??), we will not reject the hypothesis that that alleles for the MN blood type gene are in Hardy-Weinberg Equilibrium. ||




Figure 3.7.2: The area under the  $\chi_1^2$  graph that represents the p-value, the probability our hypothesis that the Locus for the MN blood group is in Hardy-Weinberg Equilibrium is correct. Since the p-value/area is so large (77 percent) we will accept our hypothesis (or Fail to Reject our hypothesis).

## 3.8 Generating Random Variables

### 3.8.1 Sampling From a Distribution

Suppose you want to do some experiments that explore the stochastic nature of a process. For example, you incorporate a random variable into a model. The first thing you must do is create some observations that have variation according to the distribution of the random variable. Most often, the method used to do this is to invert the Cumulative Distribution Function (described below) and feeding it a random number between 0 and 1.<sup>3</sup>

It is possible to convert from a uniform(0,1) random variable to something nonuniform, and also the reverse. Although this may sound backward, and in fact is, we will start by describing how to do the reverse.

Figure ?? shows a typical exponential distribution, which starts at 1 and ends near zero. The Cumulative Distribution Function (or CDF) for the same distribution can be used to map an exponential random variable to a uniform(0,1) variable and is shown in Figure ?. The CDF for any distribution is defined as the cumulative probability from the smallest number in the domain to a specific point in the dis-

<sup>3</sup>Most computer languages have standard routines that do this. For example, `rand()` in Perl and in C there is `rand()` and `random()`, which both return random numbers between 0 and `RAND_MAX`




Figure 3.8.1: PDF of an exponential process with rate = 1:  $pdf(x) = e^{-x}$

tribution,  $p$ . That is,  $P(X \leq p)$ . If  $X$  is distributed by the exponential function  $f(x) = e^{-x}$ , then

$$\begin{aligned} P(X \leq p) &= \int_0^p e^{-x} dx \\ &= [-e^{-x}]_0^p \\ &= 1 - e^{-p}. \end{aligned}$$

Since the total area under any distribution is always 1, that is for any distribution,  $g(x)$ ,

$$\int_{-\infty}^{\infty} g(x) dx = 1$$

it is clear that the CDF for any distribution can be used to map any random variable from that distribution to a number between 0 and 1. Since no part in the CDF will be visited any more than any other part, the mapping is to a uniform(0,1) distribution.

To map a uniform(0,1) random variable to an exponential distribution you simply do the reverse. That is, use the inverse of the exponential distribution CDF,  $x = -\log(1-y)$ , and apply it to a uniform(0,1) variable.

Here is some Octave code that demonstrates this transformation (see Figure ?? for the histogram):

```
octave:1> uniforms = rand(50, 1);
octave:2> exps = -log( 1 - uniforms);
octave:3> hist(exps) # plot histogram
```

This method, in general works very well for the subset of distributions that are formed by transforming uniform random variables. These include the Beta, Gamma, Chi-squared, F, Exponential, Double Exponential and Weibull distributions. Other distributions rely on tricks to generate random variables.

Despite the fact that this text attempts to avoid describing tricks at all costs, due to their utility we will list two of them.

### 3.8.2 Generating Normally Distributed Random Variables

There are two primary methods for generating normally distributed random variables. The first method relies on the central limit theorem which states that if  $EX = \mu$  and  $\text{Var}(X) = \sigma^2$ , then for independent and identically distributed samples from any distribution  $\bar{X}$  has an approximate  $N(\mu, \sigma^2/n)$  distribution, where  $n$  is the sample size. The utility in this first method is that it is very easy to remember off the top of your head and is relatively easy to compute with a computer.

The second method uses a direct transformation, and, while being just as easy to compute using a computer, is a little tricky to remember. This method is called the Box-Muller algorithm. The steps involved are:

1. Generate two independent uniform(0,1) variables,  $U_1, U_2$ .
2. Let  $R = \sqrt{-2 \log(U_1)}$  and  $\theta = 2\pi U_2$
3. Let  $X = R \cos(\theta)$  and  $Y = R \sin(\theta)$ .

where  $X$  and  $Y$  are independent normal(0,1) random variables. See Appendix ?? for an Octave program that implements this algorithm.

### 3.8.3 Generating Random Variables From a Discrete Distribution

For this, let us consider the following hypothetical discrete distribution:

$$f(x) = \begin{cases} 0 & x < 0 \\ 0.25 & 0 < x < 1 \\ 0.5 & 1 < x < 2 \\ 0.25 & 2 < x < 3 \\ 0 & 3 < x \end{cases} \quad (3.8.1)$$

Thus, with this distribution, you would expect a quarter of the data points in a random sample to come from between 0 and 1, half to come from between 1 and 2 and the remaining quarter to come from between 2 and 3.

In order to simulate a random sample from this distribution, we can map it to an interval ranging between 0 and 1. This can be done by simply taking




Figure 3.8.2: CDF of an exponential process with rate = 1:  $\text{cdf}(x) = 1 - e^{-x}$




Figure 3.8.3: A histogram of 50 exp(1) random variables




Figure 3.8.4: A histogram of 200 simulated random variables from `generate_discretet.m`.

length of each interval, multiplying it by its probability and storing the new value in an array. For this distribution, the first index of the array would contain 0.25, the second 0.5 and the third 0.25. To generate the random sample, select a random number between 0 and 1 and check if it is greater than the value in the first index. If it is not, then the point falls between 0 and 1 in our distribution. We can then choose another random number between 0 and 2 to determine which particular value it is. If the first random number is greater than the value at the first index, then we check to see if it is greater than the values at the first and second indices combined ( $0.25 + 0.5$ ) and continue this procedure until we have established which segment in the distribution it comes from and then determine the exact value from that segment by choosing another random number. This particular example has been coded in octave and the program can be found in Appendix ???. A histogram of the output can be seen in Figure ?? where the sample size was set to 200.

## 3.9 Parameter Estimation Using Maximum Likelihood

### 3.9.1 Overview

Maximum Likelihood is a method for estimating parameters for distributions. For example, if you have a set of independent<sup>4</sup> data points,  $\mathbf{X}$ , where  $\mathbf{X}$  is a vec-

<sup>4</sup>Independent simply means that knowing the value of one specific data point does not tell you anything about the value of any of the other data points. For example, if our data consisted of the results of tossing a coin, knowing that  $x_i$  landed heads

tor of points that have been collected from the same normal distribution where  $\mu$  and  $\sigma^2$  are unknown, then maximum likelihood would derive  $\hat{\mu}$  and  $\hat{\sigma}^2$ , estimates for  $\mu$  and  $\sigma^2$  respectively. The maximum likelihood estimates (MLEs) are such that they maximize the probability or likelihood of the data,  $\mathbf{X}$ .

### 3.9.2 Method

Maximum Likelihood simply uses all those Max/Min strategies that we learned in high-school calculus and then promptly forgot.

Here's the general strategy in for solving for the value of a parameter that maximizes the probability of the data:

1. Take the first derivative of the function with respect to the parameter that you want to solve for.
2. Set the derivative equal to zero and attempt to solve for the parameter.
3. If you come up with a single solution, take the second derivative of the original equation with respect to the parameter, substitute in your solution for the parameter and then check to see that it is less than zero. If so, then you have found the value that maximizes the function. (This has worked in almost every situation I have encountered.)
4. If you come up with multiple solutions, check all the solutions and check the endpoints of the range as well. (You almost never have to do this.)

Often times the log of the likelihood function is maximized instead of just the likelihood function. This is because it is almost easier to work with the log of the likelihood function than the likelihood function itself. We can justify this simplification because all probability distributions are non-negative for the domain of  $x$ , and the function  $\log[x]$  is an increasing function in  $x$ , thus, the solution for the parameter that maximizes the probability distribution given the data is the same as the maximum of the natural logarithm of the distribution given the data. Also, we'll use the notation,  $\mathcal{L}(\theta|\mathbf{X})$  to mean *The maximum with respect to  $\theta$  (the parameter that we want to estimate) of the probability of the data,  $\mathbf{X}$* . It is also worth noting that most statisticians use "log" to mean "natural log" or "ln".

would not tell us a thing about whether  $x_j$  landed heads or tails.

**Example 3.9.2.1**

From the overview, let's assume that we have  $\mathbf{X}$ , a vector of  $n$  independent data points,  $x_1, \dots, x_n$ , collected from the same normal distribution where both  $\mu$  and  $\sigma^2$  are unknown. Since each element in  $\mathbf{X}$  is independent, the probability of the data as a whole is the product of the probability of each element in  $\mathbf{X}$ .<sup>5</sup>

We will begin by finding an estimate for  $\mu$ . To do this we will assume that we know  $\sigma^2$ .

$$\begin{aligned} \mathcal{L}(\mu, \sigma^2 | \mathbf{X}) &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(x_i - \mu)^2} \\ &= \frac{1}{(2\pi\sigma^2)^{n/2}} e^{-\frac{1}{2\sigma^2} \sum (x_i - \mu)^2}, \end{aligned}$$

and

$$\begin{aligned} \log [\mathcal{L}(\mu, \sigma^2 | \mathbf{X})] &= \frac{-n}{2} \log[2\pi] - \frac{n}{2} \log[\sigma^2] \\ &\quad - \frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2. \end{aligned} \quad (3.9.1)$$

The partial derivative with respect to  $\mu$  is,

$$\frac{\partial}{\partial \mu} \log [\mathcal{L}(\mu, \sigma^2 | \mathbf{X})] = \frac{1}{\sigma^2} \sum_{i=1}^n (x_i - \mu) \stackrel{\text{set}}{=} 0,$$

Thus,

$$\begin{aligned} \frac{1}{\sigma^2} \sum_{i=1}^n (x_i - \mu) &= 0 \\ \sum_{i=1}^n (x_i - \mu) &= 0 \\ \sum_{i=1}^n x_i - n\mu &= 0 \\ n\mu &= \sum_{i=1}^n x_i \\ \hat{\mu} &= \frac{1}{n} \sum_{i=1}^n x_i = \bar{X}. \end{aligned} \quad (3.9.2)$$

<sup>5</sup>For example, if your data set was two heads when a coin is tossed twice, then the probability of the data is  $(1/2)(1/2) = 1/4$ , since the probability of getting heads on any one toss is  $1/2$ .

Verifying that  $\hat{\mu}$  is indeed a maximum requires us to take the second derivative of Equation ?? and make sure it is negative.

$$\frac{\partial^2}{(\partial \mu)^2} \log [\mathcal{L}(\mu, \sigma^2 | \mathbf{X})] = \frac{-n}{\sigma^2} < 0.$$

Thus, since  $\hat{\mu}$  is the only extreme point, is indeed a maximum.

Now we will solve for  $\hat{\sigma}^2$ , the MLE of  $\sigma^2$ . Starting from Equation ?? and substituting in our solution for  $\mu$ , we can take the partial derivative with respect to  $\sigma^2$ . Thus,

$$\frac{\partial}{\partial \sigma^2} \log [\mathcal{L}(\sigma^2 | \bar{X}, \mathbf{X})] = \frac{-n}{2\sigma^2} + \frac{1}{2\sigma^4} \sum_{i=1}^n (x_i - \bar{X})^2 \stackrel{\text{set}}{=} 0,$$

and

$$\begin{aligned} \frac{-n}{2\sigma^2} + \frac{1}{2\sigma^4} \sum_{i=1}^n (x_i - \bar{X})^2 &= 0 \\ -n\sigma^2 + \sum_{i=1}^n (x_i - \bar{X})^2 &= 0 \\ \sigma^2 &= \frac{1}{n} \sum_{i=1}^n (x_i - \bar{X})^2. \end{aligned} \quad (3.9.3)$$

To verify that our solution for  $\hat{\sigma}^2$  is indeed a maximum, we have,

$$\frac{\partial^2}{(\partial \sigma^2)^2} \log [\mathcal{L}(\sigma^2 | \bar{X}, \mathbf{X})] = \frac{n}{2\sigma^4} - \frac{1}{\sigma^6} \sum_{i=1}^n (x_i - \bar{X})^2.$$

Substituting in our solution for  $\sigma^2$  we have,

$$\begin{aligned} \frac{n^3}{2 \sum (x_i - \bar{X})^2} - \frac{n^3}{\sum (x_i - \bar{X})^6} \sum (x_i - \bar{X})^3 \\ = \frac{n^3}{2 \sum (x_i - \bar{X})^2} - \frac{n^3}{\sum (x_i - \bar{X})^2} < 0, \end{aligned}$$

and thus, our solution for  $\sigma^2$  is also a maximum.

One final note before we conclude this example. If we had attempted to solve for the MLE for  $\sigma^2$  before we solved for  $\hat{\mu}$ , then we would have ended up with the solution

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2$$

which still contains the unknown parameter  $\mu$ . At this point, we would have to pause in our derivation of  $\hat{\sigma}^2$  and solve for  $\hat{\mu}$ . Once we had a solution for  $\hat{\mu}$ , we would then substitute it in for  $\mu$  to complete our derivation of  $\hat{\sigma}^2$ . ||

### 3.9.3 Properties of Maximum Likelihood Estimates

Maximum likelihood estimates have a wealth of useful properties. I'll just list a few:

- They are *consistent*, meaning that for large  $n$ , they converge to the parameters that they estimate.
- They are *asymptotically efficient*, that is, for large  $n$ , they have minimal variance.

## 3.10 Likelihood Ratio Tests

### 3.10.1 General Overview

Likelihood ratio tests are ratios of distributions using parameters derived using both constrained and unconstrained maximum likelihood. That is, the likelihood ratio test,  $\lambda$  is,

$$\lambda = \frac{\text{constrained maximum likelihood}}{\text{unconstrained maximum likelihood}}, \quad (3.10.1)$$

where the constraint placed on the MLEs in the numerator is the hypothesis that you want to test. In Section ?? we saw how to solve for unconstrained MLEs. In the following examples we will see how to solve for and work with constrained MLEs.

The closer the ratio in Equation ?? is to 1, the more probable that the hypothesis that we are testing is true. The closer this ratio is to 0, the less likely that the hypothesis is correct. Almost all statistical tests can be derived from likelihood ratio tests. As usual, the best way to get a grasp of this concept is to see a few examples.

#### Example 3.10.1.1

Imagine that we have a set of data,  $\mathbf{X}$ , as described in Example ??, and we want to test to see if  $\mu = 3$ . That is, let the null hypothesis be  $H_0 : \mu = 3$ . In Example ?? we derived the unconstrained maximum likelihood estimates for  $\mu$  and  $\sigma^2$  (see Equations ?? and ??). In this case, to derive the constrained MLEs we simply substitute in the value 3 wherever  $\hat{\mu}$  is used, including the derivation of  $\hat{\sigma}^2$ . Thus,

$$\hat{\mu}_c = 3 \text{ and } \hat{\sigma}_c^2 = \frac{1}{n} \sum_{i=1}^n (x_i - 3)^2.$$

and our likelihood ratio test is:

$$\begin{aligned} \lambda &= \frac{\max_{H_0} \mathcal{L}(\mu, \sigma^2 | \mathbf{X})}{\mathcal{L}(\mu, \sigma^2 | \mathbf{X})} = \frac{\Pr(\mathbf{X} | 3, \hat{\sigma}_c^2)}{\Pr(\mathbf{X} | \hat{\mu}, \hat{\sigma}^2)} \\ &= \frac{\prod_{i=1}^n \frac{1}{\sqrt{2\pi\hat{\sigma}_c^2}} e^{-\frac{1}{2\hat{\sigma}_c^2}(x_i-3)^2}}{\prod_{i=1}^n \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} e^{-\frac{1}{2\hat{\sigma}^2}(x_i-\hat{\mu})^2}} \\ &= \frac{(2\pi\hat{\sigma}_c^2)^{-n/2} \exp\left\{-\frac{n}{2\sum(x_i-3)^2} \sum(x_i-3)^2\right\}}{(2\pi\hat{\sigma}^2)^{-n/2} \exp\left\{-\frac{n}{2\sum(x_i-\hat{\mu})^2} \sum(x_i-\hat{\mu})^2\right\}} \\ &= \frac{(\hat{\sigma}_c^2)^{-n/2} e^{\frac{n}{2}}}{(\hat{\sigma}^2)^{-n/2} e^{\frac{n}{2}}} \\ &= \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_c^2}\right)^{n/2} \end{aligned}$$

||

#### Example 3.10.1.2

Imagine that we have the same set up as we had in Example ??, only this time, the hypothesis that we want to test is  $\mu \leq 3$ . In this case, when  $\hat{\mu} \leq 3$ , we let  $\hat{\mu}_c = \hat{\mu}$ . However, when  $\hat{\mu} > 3$ , then  $\hat{\mu}_c = 3$ . Thus,

$$\lambda = \begin{cases} \frac{\Pr(\mathbf{X} | \hat{\mu}, \hat{\sigma}^2)}{\Pr(\mathbf{X} | \hat{\mu}, \hat{\sigma}^2)} = 1, & \hat{\mu} \leq 3 \\ \frac{\Pr(\mathbf{X} | 3, \hat{\sigma}^2)}{\Pr(\mathbf{X} | \hat{\mu}, \hat{\sigma}^2)} = \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_c^2}\right)^{n/2}, & \hat{\mu} > 3 \end{cases}$$

Notice that the LRT for  $H_0 : \mu \neq 3$  is the same as the LRT for  $H_0 : \mu \leq 3$  when  $\hat{\mu} > 3$ . ||

## 3.11 Solving Constrained Optimization Problems with Lagrange Multipliers

### 3.11.1 General Overview

It is often the case that you have some function, for example  $f$ , and you want to find its extreme values (maximum and or minimum). If  $f$  is a function of a single variable,  $f(x)$ , then all that's needed is to take its derivative and check where it equals zero. If we then went a little further and wanted the maximum and minimum values of  $f(x)$  when  $x$  also falls on a specific line (or is otherwise constrained by some other function  $g(x)$ ), then all we would need to do is find out where (and if) the two lines intersected and




Figure 3.11.1: A 3-D graph of two surfaces intersecting

check to see which values for  $x$  give us the largest and smallest values from  $f(x)$ .

When functions of multiple variables are involved and we are still trying to find constrained maximums and minimums, we do a similar thing. In this case, instead of picturing lines intersecting, we can imagine surfaces intersecting (see Figure ??). Regardless of how we visualize the problem, there are two primary strategies for solving this sort of problem. The first one, algebraic substitution, works well when there are only two variables involved (otherwise you run into the problem of having more variables to solve for than you have equations). In this case you simply use the constraint equation to solve for one variable in terms of another, substitute this solution into the equation you wish to find the maximum/minimum of and then take the derivative to solve for the extreme points. The second method, Lagrange multipliers, (the one we'll be covering here), is a little more fancy, but works well when there are more than two variables involved.<sup>6 7</sup>

To find the extreme values of  $f(x, y, z)$  subject to the constraint  $g(x, y, z) = k$ :

1. Find all values of  $x, y, z$  and  $\lambda$  such that

$$\nabla f(x, y, z) = \lambda \nabla g(x, y, z) \quad (3.11.1)$$

and  $g(x, y, z) = k$

<sup>6</sup>There is a third alternative, called "variational method" which is interesting, but I don't quite fully understand well enough to write about at this time

<sup>7</sup>For now we will discuss how the method works for functions of three variables, but it works fine on functions with two variables (as you'll see in Example ??) and trivially extends to functions with more.

2. Evaluate  $f(x, y, z)$  at all points,  $(x, y, z)$ , that result from the previous step.

### Example 3.11.1.1

In this example we will use equations with only two variables and demonstrate how both algebraic substitution and Lagrange multipliers result in identical solutions.

Let's find the extreme values of

$$f(x, y) = x^2 + 2y^2 \quad (3.11.2)$$

that are also on the circle

$$x^2 + y^2 = 1 \quad (3.11.3)$$

(thus, in this case,  $g(x, y) = x^2 + y^2$ ).

**Using Algebraic Substitution:** Using the constraint equation,  $x^2 + y^2 = 1$ , we'll solve for  $y^2$  by moving  $x^2$  to the other side. Thus,  $y^2 = 1 - x^2$ . We can then substitute  $1 - x^2$  for  $y^2$  into  $x^2 + 2y^2$  (Equation ??), giving us

$$\begin{aligned} f(x) &= x^2 + 2(1 - x^2) \\ &= x^2 + 2 - 2x^2 \\ &= 2 - x^2. \end{aligned} \quad (3.11.4)$$

We can now take the derivative of Equation ?? and solve for the extreme points. That is,

$$\begin{aligned} f'(x) &= -2x \\ -2x &\stackrel{\text{set}}{=} 0 \\ x &= 0. \end{aligned} \quad (3.11.5)$$

Combining Equation ?? with Equation ??, we can determine that when  $x = 0$ ,  $y = \pm 1$ .

Now we must do the same thing only this time, we must use Equation ?? to solve for  $x^2$ . That is, from Equation ??,  $x^2 = 1 - y^2$ . Substituting this into Equation ??, we then follow the same steps as before to determine that when  $y = 0$ ,  $x = \pm 1$ . We can then plug these points into Equation ?? to find that Thus  $(0, \pm 1)$  are both maximums and  $(\pm 1, 0)$  are both minimums.

**Using Lagrange Multipliers:** To use the Lagrange multiplier method, we'll first set up the system of equations defined by  $\nabla f(x, y, z) = \lambda \nabla g(x, y, z)$ :

$$2x = \lambda 2x \quad (3.11.6)$$

$$4y = \lambda 2y. \quad (3.11.7)$$



From Equation ?? we can determine that either  $x = 0$  or  $\lambda = 1$ . If  $x = 0$ , the constraint equation, (Equation ??), forces  $y = \pm 1$ . If  $\lambda = 1$ , then we can use Equation ?? to determine that  $y = 0$  and thus, from Equation ??,  $x = \pm 1$ . The points of intersection are therefore  $(0, \pm 1)$  and  $(\pm 1, 0)$ , the same solutions derived using algebraic substitution. ||

### Example 3.11.1.2

We want to find the maximum value of

$$f(x, y, z) = xyz \quad (3.11.8)$$

subject to the constraint

$$2xz + 2zy + xy = 12 \quad (3.11.9)$$

(thus,  $g(x, y, z) = 2xz + 2zy + xy$ ). In this example, the equations involve more than two variables so algebraic substitution is not an option. Thus, it makes sense to use the Lagrange multiplier method here. We will first set up the system of equations defined by  $\nabla f(x, y, z) = \lambda \nabla g(x, y, z)$ :

$$yz = \lambda(2z + y) \quad (3.11.10)$$

$$xz = \lambda(2z + x) \quad (3.11.11)$$

$$xy = \lambda(2x + 2y). \quad (3.11.12)$$

If we multiply Equation ?? by  $x$ , Equation ?? by  $y$  and Equation ?? by  $z$ , we have,

$$xyz = \lambda(2xz + xy) \quad (3.11.13)$$

$$xyz = \lambda(2yz + xy) \quad (3.11.14)$$

$$xyz = \lambda(2xz + 2yz). \quad (3.11.15)$$

We can now use these equations to solve for  $x$  and  $y$  in terms of  $z$ . Using Equations ?? and ?? we get,

$$\lambda(2xz + xy) = \lambda(2yz + xy)$$

$$2xz + xy = 2yz + xy$$

$$x = y. \quad (3.11.16)$$

Using Equations ?? and ?? we get  $y$  in terms of  $z$ , that is,

$$\lambda(2yz + xy) = \lambda(2xz + 2yz)$$

$$2yz + xy = 2xz + 2yz$$

$$y = 2z. \quad (3.11.17)$$

We can now substitute  $2z$  for both  $x$  and  $y$  into Equation ?? to solve for  $z$ :

$$2(2z)z + 2(2z)z + (2z)(2z) = 12$$

$$4z^2 + 4z^2 + 4z^2 = 12$$

$$z = 1. \quad (3.11.18)$$

Using Equations ??, ?? and ??, we can determine that the only place where both functions intersect is when  $x = 2$ ,  $y = 2$  and  $z = 1$ . Thus  $f(2, 2, 1) = 4$  is both the maximum and minimum from this function subject to the constraint, Equation ??.

## 3.12 Matrix Calculus

Basically derivatives and integrals of matrices parallel derivatives and integrals of ordinary functions. The easiest way to figure out what the derivative or integral will be is to expand the matrix and then take an element by element derivative or integral.

If we let  $\mathbf{x}$  be an  $n \times 1$  vector and let  $\mathbf{y} = f(\mathbf{x})$ , where  $\mathbf{y}$  is an  $m \times 1$  vector (for example, if  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ , where  $\mathbf{A}$  is an  $m \times n$  matrix, then  $\mathbf{y}$  will be an  $m \times 1$  vector), then

$$\frac{\partial \mathbf{y}}{\partial \mathbf{x}} = \begin{bmatrix} \frac{\partial y_1}{\partial x_1} & \dots & \frac{\partial y_m}{\partial x_1} \\ \frac{\partial y_1}{\partial x_2} & \dots & \frac{\partial y_m}{\partial x_2} \\ \vdots & \ddots & \vdots \\ \frac{\partial y_1}{\partial x_n} & \dots & \frac{\partial y_m}{\partial x_n} \end{bmatrix}. \quad (3.12.1)$$

We will also include the following to our definition:

$$\frac{\partial \mathbf{y}}{\partial \mathbf{x}'} = \left( \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \right)' \quad (3.12.2)$$

and

$$\frac{\partial \mathbf{y}'}{\partial \mathbf{x}} = \frac{\partial \mathbf{y}}{\partial \mathbf{x}}. \quad (3.12.3)$$

Now we'll list two very useful results.<sup>8</sup> If  $\mathbf{x}$  is an  $n \times 1$  vector and  $\mathbf{A}$  is an  $m \times n$  matrix of elements that are not functions of  $\mathbf{x}$ , then

$$\frac{\partial \mathbf{A}\mathbf{x}}{\partial \mathbf{x}} = \mathbf{A}'. \quad (3.12.4)$$

If  $\mathbf{A}$  is an  $n \times n$  matrix of elements that are not functions of  $\mathbf{x}$ , then

$$\frac{\partial f(\mathbf{x}'\mathbf{A}\mathbf{x})}{\partial \mathbf{x}} = (\mathbf{A} + \mathbf{A}')\mathbf{x} \quad (3.12.5)$$

and if  $\mathbf{A}$  is symmetric, that is  $\mathbf{A} = \mathbf{A}'$ , then

$$(\mathbf{A} + \mathbf{A}')\mathbf{x} = 2\mathbf{A}\mathbf{x}.$$

<sup>8</sup>Complete derivations of these results can be found in Appendix ?? and ??.

## 3.13 Linear Models

### 3.13.1 General Overview

Let's say that you are studying a type of chicken and you have reason to believe that its weight will give you some indication of how much food it will eat in a year (a fairly reasonable thing to suspect). Ideally we would like to eventually have some sort of function that we could use hen weight for input and the result would be an estimate of how much feed we might expect it to consume.

So, you go out and weigh a hen and it turns out to weigh 4.6 units and consumes 87.1 units. From this single data point, it would be impossible to tell if a hen that weighed more would eat more (which would be what we suspected) or would eat less. Thus, we go out and collect another data point. This time the hen weighs 5.1 units and eats 93.1 units. If we assumed that there was some sort of linear relationship between the hen's weight and the amount of feed it consumes, then we could use the two data points to solve for the unknown parameters in our model, using them to solve for an intercept (which we will call  $\beta_0$ ) a slope (which we will call  $\beta_1$ ). Thus, using the following two equations

$$87.1 = \beta_0 + \beta_1(4.6)$$

$$93.1 = \beta_0 + \beta_1(5.1)$$

and standard algebraic techniques, we can determine that  $\beta_0 = 31.9$  and  $\beta_1 = 12$ . Thus our model is:

$$f(x) = 31.9 + 12x. \quad (3.13.1)$$

After measuring several more points (Table ??) you realize that none of them, except for the first two which were used to create the model, fall on the line defined by  $f(x)$  (See Figure ??).

At this point we might realize that it was fairly arbitrary to decide to use the first two points to create our model. We could have used the second and the third or the fourth and fifth, but using any specific pair of points to define our model doesn't make it any less arbitrary. What we would really like to do is use all of the data that we have collected to create our model. Since it is obvious that all of the data does not fall on a single line<sup>9</sup> we would like to create

<sup>9</sup>It is important to note, that just because the data does not all fall on a single line, doesn't mean that the model is not linear. There could have been errors in measurement, both human and mechanical, that cause the data to deviate from a line.

Body weight, $\bar{X}$	Food Consumption, $\bar{Y}$
4.6	87.1
5.1	93.1
4.8	89.8
4.4	91.4
5.9	99.5
4.7	92.1
5.1	95.5
5.2	99.3
4.9	93.4
5.1	94.4

Table 3.13.1: Average body weight  $X$  and food consumption  $Y$  for 50 hens from each of 10 White Leghorn strains (350-day period). *Source:* Plagiarized from Steel, Torrie and Dickey [?]. Data from S. C. King, Purdue University




Figure 3.13.1: A plot of the Leghorn data from Table ?? with a line drawn using the first two points to define the slope and the intercept (Equation ??). Notice how poorly this line estimates the other data points. For example, with a single exception, the estimates made by Equation ?? are low. Compare this with the graph shown in Figure ??.




Figure 3.13.2: A plot of the Leghorn data from Table ?? with a line drawn using Least Squares to estimate the slope and the intercept. Notice how even though this line passes through fewer points than Equation ??, shown in Figure ??, it tends to closer to the majority of the data.

our model in such a way that the difference between the points that the model predicts and the observed data is minimized (see Figure ??). This section concerns itself with describing both a method for creating models that achieve this called *Least Squares*, and a means to evaluate the the properties of these models. This method works well with a wide range of data (not just simple  $(x, y)$  pairs) and this will be seen in the examples.<sup>10</sup>

Least squares is a method for estimating parameters for linear functions (or, in more technical jargon, functions that are linear with respect to its coefficients<sup>11</sup>) such that they minimize the sum of squares of differences between the  $y$ -values of the data points and the corresponding  $y$ -values of the approximating function.

We start by considering a linear model of the form

$$y_i = \beta_0 + \beta_1 x_{i,1} + \dots + \beta_m x_{i,m} + \epsilon_i, \quad (3.13.2)$$

where  $i = 1, \dots, n$  is the number of observations. This system of  $n$  equations can be written in matrix

<sup>10</sup>However, see Example ?? for the solution to this current conundrum!

<sup>11</sup>For a function,  $f(x)$ , to be considered linear with respect to its coefficients means that if the function were considered to be a function of the coefficients,  $g(\mathbf{c})$ , then  $g(\alpha\mathbf{c}) = \alpha g(\mathbf{c})$ . For example, the function,  $f(x) = c_0 + c_1 x$ , can be written in terms of  $\mathbf{c}$ ,  $g(\mathbf{c}) = c_0 + c_1 x$  and  $g(\alpha\mathbf{c}) = \alpha c_0 + \alpha c_1 x = \alpha(c_0 + c_1 x) = \alpha f(\mathbf{c})$ . Another example of a function that is linear with respect to its coefficients is  $f(x) = c_0 \sin(x) + c_1 e^x$ , because  $g(\alpha\mathbf{c}) = \alpha g(\mathbf{c})$ . An example of a function that is not linear with respect to its coefficients is  $f(x) = \sin(c_0 x) + c_1 x^2$ , since  $g(\alpha\mathbf{c}) = \sin(\alpha c_0 x) + \alpha c_1 x^2 \neq \alpha(\sin(c_0 x) + c_1 x^2) = \alpha g(\mathbf{c})$

notation quite concisely with,

$$\mathbf{Y} = \mathbf{X}\beta + \epsilon, \quad (3.13.3)$$

where  $\mathbf{Y}$ , called the dependent variable, is an  $n \times 1$  vector of observed measurements<sup>12</sup>,  $\beta$  is an  $1 \times m$  vector of unknown model parameters,  $\mathbf{X}$ , called the independent variables or the *design matrix*, is an  $n \times m$  matrix of independent variable values and  $\epsilon$  is the measurement noise.<sup>13</sup>

### 3.13.2 Setting up $\mathbf{Y}$ and $\mathbf{X}$

Setting  $\mathbf{Y}$  is always a strait forward procedure: you simply fill the vector with the measured values. Setting up  $\mathbf{X}$ , the design matrix, however, depends on the type of data you have as well as the model you are trying to fit. All of this is best explained with a series of examples.

#### Example 3.13.2.1

First we will show how to set up the design matrix when there is a single independent variable involved.

If we are given the data,

$$\begin{array}{c|c} y & x \\ \hline 1 & 0 \\ 4 & 1 \\ 5 & 2 \end{array},$$

and the function we wish to fit using least squares is,

$$y = \beta_0 + \beta_1 x,$$

then

$$\mathbf{Y} = \begin{bmatrix} 1 \\ 4 \\ 5 \end{bmatrix} \text{ and } \mathbf{X} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{bmatrix}.$$

The column of 1s in  $\mathbf{X}$  represents  $\beta_0 \times 1 = \beta_0$ .

If the equation we wanted to fit was quadratic,

$$y = \beta_0 + \beta_1 x + \beta_2 x^2,$$

then

$$\mathbf{Y} = \begin{bmatrix} 1 \\ 4 \\ 5 \end{bmatrix} \text{ and } \mathbf{X} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 4 \end{bmatrix}.$$

<sup>12</sup>In Section ??, *Linear Models with Multiple Dependent Variables* we will generalize the test developed here for multiple dependent variables.

<sup>13</sup>I like to use this concept instead of calling  $\epsilon$  "error" which it is not. It simply reflects the limits of our ability to capture the totality of what is going on. With perfect models, we'd be able to capture the thermal noise generated by molecular motion and have a perfect fit. So - errors - NO, noise - YES.

If the equation was

$$\beta_0 \sin(x) + \beta_1 x^2,$$

then

$$\mathbf{Y} = \begin{bmatrix} 1 \\ 4 \\ 5 \end{bmatrix} \text{ and } \mathbf{X} = \begin{bmatrix} \sin(0) & 0 \\ \sin(1) & 1 \\ \sin(2) & 4 \end{bmatrix}.$$

In general, for the table of data,

$y$	$x$
$y_1$	$x_1$
$y_2$	$x_2$
$\vdots$	$\vdots$
$y_m$	$x_m$

and any function that is linear with respect to the coefficients,

$$y = \beta_0 f_0(x) + \beta_1 f_1(x) + \cdots + \beta_n f_n(x),$$

then

$$\mathbf{X} = \begin{bmatrix} f_0(x_0) & f_1(x_0) & \cdots & f_n(x_0) \\ f_0(x_1) & f_1(x_1) & \cdots & f_n(x_1) \\ \vdots & \vdots & \ddots & \vdots \\ f_0(x_m) & f_1(x_m) & \cdots & f_n(x_m) \end{bmatrix}.$$

||

### Example 3.13.2.2

If we are given a dataset that contains multiple independent variables, for example:

$y$	$x_1$	$x_2$	$x_3$
1	0	4.3	0.2300
4	1	2.3	0.0010
5	2	7.5	0.0004

and we want to find a fit for the function

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3,$$

then you would end up with

$$\mathbf{Y} = \begin{bmatrix} 1 \\ 4 \\ 5 \end{bmatrix} \text{ and } \mathbf{X} = \begin{bmatrix} 1 & 0 & 4.3 & 0.2300 \\ 1 & 1 & 2.3 & 0.0010 \\ 1 & 2 & 7.5 & 0.0004 \end{bmatrix}.$$

||

### Example 3.13.2.3

Sometimes the independent variable is a list of treatments and the dependent variable consists of a list of values measured after each treatment. For example, if you have the data set,

Treatment 1	Treatment 2	Treatment 3
13.2	10.4	15.2
12.8	9.7	15.0
13.5	10.2	15.8

we can still use a linear model,

$$y = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3$$

and estimate the parameters  $\beta_1$ ,  $\beta_2$  and  $\beta_3$ . However, in this case,  $x_i$  consists of a 0 or a 1, depending on which treatment a given  $y$  value was collected from. Thus,

$$\mathbf{Y} = \begin{bmatrix} 13.2 \\ 12.8 \\ 13.5 \\ 10.4 \\ 9.7 \\ 10.2 \\ 15.2 \\ 15.0 \\ 15.8 \end{bmatrix}, \quad \mathbf{X} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

For a discussion of alternative design matrices (some of which are used, for historical reasons, more often than this one) for this type of data set, see Appendix ??.

### 3.13.3 Parameter Estimation: The Least Squares Method

Given data for the dependent and independent variables,  $\mathbf{X}$  and  $\mathbf{Y}$ , how should we estimate the values for  $\beta$ , the model parameters? For this we can use the least squares procedure. That is, estimate  $\beta$  by minimizing the total squared differences between observed and predicted values. The difference between the observed and predicted values, often times called the *residual*, is, in matrix notation,  $\mathbf{Y} - \mathbf{X}\beta$ . The squared residual is  $(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta)$ . Thus,

$$\begin{aligned} \mathbf{F}(\beta) &= (\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta) \\ &= \mathbf{Y}'\mathbf{Y} - \mathbf{Y}'\mathbf{X}\beta - \beta'\mathbf{X}'\mathbf{Y} + \beta'\mathbf{X}'\mathbf{X}\beta. \end{aligned} \tag{3.13.4}$$

To minimize Equation ??, we take its derivative with respect to  $\beta$ , set it equal to zero and solve for

$\beta$ .

$$\begin{aligned} & \frac{\partial}{\partial \beta} \mathbf{Y}'\mathbf{Y} - \mathbf{Y}'\mathbf{X}\beta - \beta'\mathbf{X}'\mathbf{Y} + \beta'\mathbf{X}'\mathbf{X}\beta \\ &= -\mathbf{X}'\mathbf{Y} - \mathbf{X}'\mathbf{Y} + 2\mathbf{X}'\mathbf{X}\beta \\ &= -2\mathbf{X}'\mathbf{Y} + 2\mathbf{X}'\mathbf{X}\beta \stackrel{\text{set}}{=} \mathbf{0}, \end{aligned}$$

and thus<sup>14</sup>

$$\hat{\beta} = [\mathbf{X}'\mathbf{X}]^{-1}\mathbf{X}'\mathbf{Y}. \quad (3.13.5)$$

If we substitute our estimated parameters,  $\hat{\beta}$ , into Equation ??, we get the following simplification for calculating the squared residual:

$$\begin{aligned} \mathbf{F}(\hat{\beta}) &= \mathbf{Y}'\mathbf{Y} - \mathbf{Y}'\mathbf{X}\hat{\beta} - \hat{\beta}'\mathbf{X}'\mathbf{Y} + \hat{\beta}'\mathbf{X}'\mathbf{X}\hat{\beta} \\ &= \mathbf{Y}'\mathbf{Y} - \mathbf{Y}'\mathbf{X}\hat{\beta} - \hat{\beta}'\mathbf{X}'\mathbf{Y} + \hat{\beta}'\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} \\ &= \mathbf{Y}'\mathbf{Y} - \mathbf{Y}'\mathbf{X}\hat{\beta} - \hat{\beta}'\mathbf{X}'\mathbf{Y} + \hat{\beta}'\mathbf{X}'\mathbf{Y} \\ &= \mathbf{Y}'\mathbf{Y} - \mathbf{Y}'\mathbf{X}\hat{\beta}. \end{aligned} \quad (3.13.6)$$

### 3.13.4 Properties of $\hat{\beta}$

If we assume that the elements in the noise vector,  $\epsilon$ , are independent and normally distributed<sup>15</sup> random variables with  $\mu_\epsilon = 0$  and  $\sigma_\epsilon^2 = \sigma^2$ , (which is not terribly unreasonable to do since noise can come from all kinds of sources and once we add them all up, the central limit theorem kicks into effect,) then we can determine if  $\hat{\beta}$  is biased and what its variance is.

Before we start, however, we will note that the assumption that  $\epsilon$  are independent and identically distributed normal( $0, \sigma^2$ ) variables implies that  $\mathbf{Y}$  is also normally distributed with mean  $\mathbf{X}\beta$  and variance  $\sigma^2$ . This is because  $\mathbf{Y} = \mathbf{X}\beta + \epsilon$  and  $\mathbf{X}\beta$  functions as a location parameter.

<sup>14</sup>We can easily verify that this solution for  $\hat{\beta}$  is a minimum by taking the second derivative of Equation ?? with respect to  $\beta$  and observing that when  $\mathbf{X}$  is not completely filled with zeros, the resulting quantity,  $2\mathbf{X}'\mathbf{X}$ , will be positive.

<sup>15</sup>It is possible to use distributions other than the normal as long as each  $\epsilon_i$  is an independent variable with mean 0 (zero) and variance  $\sigma^2$ . These conditions are called *Gauss-Markov Conditions*. However, when you use a normal distribution, the least squares estimates are the same as the maximum likelihood estimates and are thus best unbiased estimators, which is a good thing.

First, we will show that  $\hat{\beta}$  is unbiased.

$$\begin{aligned} \mathbf{E}\hat{\beta} &= \mathbf{E}\{[\mathbf{X}'\mathbf{X}]^{-1}\mathbf{X}'\mathbf{Y}\} \\ &= [\mathbf{X}'\mathbf{X}]^{-1}\mathbf{X}'(\mathbf{E}\mathbf{Y}) \\ &= [\mathbf{X}'\mathbf{X}]^{-1}\mathbf{X}'(\mathbf{X}\beta) \\ &= [\mathbf{X}'\mathbf{X}]^{-1}\mathbf{X}'\mathbf{X}\beta \\ &= \beta. \end{aligned} \quad (3.13.7)$$

Now we will derive the variance of  $\hat{\beta}$ . However, before we get into it, let me first point out that  $\mathbf{E}(\epsilon\epsilon') = \sigma^2$ . This can be easily shown using the facts that  $\mathbf{Var}(\epsilon) = \sigma^2$ ,  $\mathbf{E}\epsilon = 0$  and the definition of variance. That is,

$$\begin{aligned} \mathbf{Var}(\epsilon) &= \mathbf{E}(\epsilon\epsilon') - \mathbf{E}\epsilon\mathbf{E}\epsilon' \\ &= \mathbf{E}(\epsilon\epsilon') - 0 \\ \sigma^2 &= \mathbf{E}(\epsilon\epsilon'). \end{aligned}$$

With that little bit of extra information in hand, we are now ready to derive the variance of  $\hat{\beta}$ .

$$\begin{aligned} \mathbf{Var}(\hat{\beta}) &= \mathbf{E}\{(\hat{\beta} - \mathbf{E}\hat{\beta})(\hat{\beta} - \mathbf{E}\hat{\beta})'\} \\ &= \mathbf{E}\{[(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} - \beta][(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} - \beta]'\} \\ &= \mathbf{E}\{[(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'(\mathbf{X}\beta + \epsilon) - \beta] \\ &\quad [(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'(\mathbf{X}\beta + \epsilon) - \beta]'\}, \\ &\quad \text{since } \mathbf{Y} = \mathbf{X}\beta + \epsilon \\ &= \mathbf{E}\{[\beta + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\epsilon - \beta] \\ &\quad [\beta + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\epsilon - \beta]'\} \\ &= \mathbf{E}\{[(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\epsilon][(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\epsilon]'\} \\ &= \mathbf{E}\{(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\epsilon\epsilon'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\} \\ &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{E}(\epsilon\epsilon')\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \\ &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\sigma^2 \\ &= (\mathbf{X}'\mathbf{X})^{-1}\sigma^2. \end{aligned}$$

An alternative and shorter derivation of this same variance is as follows:

$$\begin{aligned} \mathbf{Var}(\hat{\beta}) &= \mathbf{Var}((\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}) \\ &= [(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']\mathbf{Var}(\mathbf{Y})[(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}]' \\ &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\sigma^2 \\ &= (\mathbf{X}'\mathbf{X})^{-1}\sigma^2. \end{aligned} \quad (3.13.8)$$

### 3.13.5 Hypothesis Testing

At this point we are ready to start evaluating our estimated parameters. For example, if a parameter  $\beta_i$  is very close to zero, it could mean that the data scaled by  $\beta_i$  is insignificant in our model. If all of the values for the parameters are similar to each other, it could mean that there is no difference between the different data sets that form  $\mathbf{X}$ . This is where statistics comes in. Depending on the variability in our data, our hypothesis that  $\beta_i$  is sufficiently close to zero to be insignificant, may or may not be rejected. Statistics tells us how close, based on the data, the parameter needs to be to zero for it to be close enough accept our hypothesis (or, how far does the parameter need to be from zero to reject our hypothesis).

#### THE MAIN IDEA - Hypothesis Testing:

This section contains a lot of nasty math and so, before we get there, I want to make sure that you, the reader, stay focused on the main idea. The main idea behind all of this nasty math is this: We do two different least squares estimations for the parameters for our model. In one estimation we make no assumptions about what any of those parameters might be, we simply try to determine the values that best fit the data. In the other estimation, we assume that our hypothesis about the parameters (i.e. that some of them are equal to zero, or that all of the parameters are equal to each other) is correct and require least squares to generate parameter values given these constraints. If our hypothesis turns out to be correct, these two different parameter estimations will be quite similar. If our hypothesis is not correct, then the estimations will be different. Our goal here is to develop a statistical procedure for evaluating the similarity or difference between the two estimations.

If you can understand this idea, then you may just want to focus primarily on the examples and only briefly skim over the derivation of the statistical procedure.

The first thing we need to learn how to do is to convert our hypotheses into matrix notation so that we can use them to constrain our estimates of the parameters, just as we did when we created LRTs (see Section ??). The best way to learn this is to just see a few examples.

#### Example 3.13.5.1

If we want to test if  $\beta_i$  is zero, then we set up the matrices (or, rather, vectors in this case):

$$\begin{bmatrix} 0 & \dots & 1_i & \dots & 0 \end{bmatrix} \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_i \\ \vdots \\ \beta_p \end{bmatrix} = 0$$

$$\beta_i = 0$$

to to be the constraint that we use when we try to estimate  $\beta$ . The matrix on the left is generally called the *contrast matrix*, and we will refer to it with  $\mathbf{C}$ . ||

#### Example 3.13.5.2

If our model had four parameters and we wanted to test to see if they were all equal<sup>16</sup>, we would use the following constraint.

$$\begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$\begin{bmatrix} \beta_1 - \beta_4 \\ \beta_2 - \beta_4 \\ \beta_3 - \beta_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

If, given the data, there is a high probability that this system of equations is correct, then all of the parameters must all be equal. ||

In general our hypotheses will take the form

$$\mathbf{C}\beta = \theta,$$

where  $\mathbf{C}$  is an  $t \times p$  matrix where  $t$  is the number of tests in our hypothesis and  $p$  is the number of parameters we are estimating and  $\theta$  is just a generalization of the vector of zeros that tend to end up right side of the constraints. It is usual for  $\theta$  to contain only zeros but sometimes it can contain other numbers as well.

We can test our hypotheses by creating a LRT. That is, by comparing the unconstrained optimization of the least squares residual with a constrained optimization of the least squares residual. In Section ?? we solved for the unconstrained optimization and the result was Equation ???. To solve for the constrained optimization, that is

$$\min_{\mathbf{C}\beta = \theta} (\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta), \quad (3.13.9)$$

<sup>16</sup>This would amount to an ANOVA test. See Example ?? for a full treatment of this.

we will use the method of Lagrange Multipliers (see Section ?? for details on this method). Thus, we want to find solutions to the equations:

$$\nabla(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta) = \lambda' \nabla \mathbf{C}\beta - \boldsymbol{\theta} \quad (3.13.10)$$

and

$$\mathbf{C}\beta - \boldsymbol{\theta} = 0,$$

or

$$\mathbf{C}\beta = \boldsymbol{\theta}. \quad (3.13.11)$$

Taking the gradient of Equation ??, we have

$$\begin{aligned} \nabla(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta) &= \lambda' \nabla(\mathbf{C}\beta - \boldsymbol{\theta}) \\ \frac{\partial}{\partial \beta}(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta) &= \frac{\partial}{\partial \beta}(\lambda' \mathbf{C}\beta - \lambda' \boldsymbol{\theta}) \\ -2\mathbf{X}'\mathbf{Y} + 2\mathbf{X}'\mathbf{X}\beta &= \mathbf{C}'\lambda \\ -\mathbf{X}'\mathbf{Y} + \mathbf{X}'\mathbf{X}\beta &= \frac{1}{2}\mathbf{C}'\lambda \\ \mathbf{X}'\mathbf{X}\beta &= \mathbf{X}'\mathbf{Y} + \frac{1}{2}\mathbf{C}'\lambda. \end{aligned}$$

Thus, in terms of  $\lambda$ , the constrained optimization for  $\beta$  is

$$\begin{aligned} \hat{\beta}_c &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} + \frac{1}{2}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}'\lambda \\ &= \hat{\beta} + \frac{1}{2}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}'\lambda. \end{aligned} \quad (3.13.12)$$

We can solve for  $\lambda$  by multiplying both sides of the equation by  $\mathbf{C}$  and incorporating the constraint, Equation ??, into our solution. That is,

$$\mathbf{C}\beta = \mathbf{C}\hat{\beta} + \frac{1}{2}\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}'\lambda = \boldsymbol{\theta},$$

from Equation ??, and,

$$\boldsymbol{\theta} - \mathbf{C}\hat{\beta} = \frac{1}{2}\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}'\lambda,$$

thus,

$$\lambda = 2[\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}']^{-1}[\boldsymbol{\theta} - \mathbf{C}\hat{\beta}]. \quad (3.13.13)$$

We can substitute this solution into Equation ?? giving us the optimal solution for  $\beta$  in terms of  $\mathbf{X}$ ,  $\mathbf{Y}$ , and  $\mathbf{C}$  only.

$$\begin{aligned} \hat{\beta}_c &= \hat{\beta} \\ &+ (\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}[\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}']^{-1}(\boldsymbol{\theta} - \mathbf{C}\hat{\beta}). \end{aligned} \quad (3.13.14)$$

The LRT for the hypothesis  $\mathbf{C}\beta = 0$  is thus,

$$\lambda = \frac{\max_{\mathbf{C}\beta - \boldsymbol{\theta} = 0} \mathcal{L}(\beta, \sigma^2 | \mathbf{Y})}{\mathcal{L}(\beta, \sigma^2 | \mathbf{Y})}.$$

The constrained and unconstrained estimates for  $\beta$  derived using least squares are equivalent to those found using maximum likelihood methods (see Section ??) when we assume that the elements of  $\epsilon$ , are independent and normally distributed (just as we did when we derived the mean and variance for  $\hat{\beta}$ ). Intuitively this makes sense because the normal distribution is unimodal and symmetric about the mean. Least squares attempts to position the mean of the distribution so that the distance from the data points and the mean is minimized, and thus, the data points will tend to be in the middle of the distribution where probability is highest. Maximum likelihood methods attempt to position the mean so that the data points have maximum probability of occurring, and thus, clustered around the mean. Analytically, this equivalence is easy to derive. Using the maximum likelihood method described in Section ??, we have:

$$\mathcal{L}(\beta | \sigma^2, \mathbf{Y}) = (2\pi\sigma^2)^{-n/2} e^{-\frac{1}{2\sigma^2}(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta)}$$

$$\log[\mathcal{L}] = \frac{-n}{2} \log[2\pi\sigma^2] - \frac{1}{2\sigma^2}(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta)$$

$$\frac{\partial}{\partial \beta} \log[\mathcal{L}] = -\frac{1}{2\sigma^2}(-2\mathbf{X}'\mathbf{Y} + 2\mathbf{X}'\mathbf{X}\beta) \stackrel{\text{set}}{=} 0$$

Solving for  $\beta$  gives us the same result we found using least squares. That is,

$$\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$$

Using maximum likelihood to estimate  $\beta$  under the constraint (Equation ??) results in the least square estimate,  $\hat{\beta}_c$  (Equation ??).

Now we need to solve for the unconstrained and constrained estimates of  $\sigma^2$ . We'll begin with solving for the unconstrained estimate,  $\hat{\sigma}^2$  using maximum likelihood methods.<sup>17</sup> That is,

$$\frac{\partial}{\partial \sigma^2} \log[\mathcal{L}] = \frac{-n}{2\sigma^2} + \frac{1}{2\sigma^4}(\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta}) \stackrel{\text{set}}{=} 0,$$

thus,

$$\begin{aligned} -n\sigma^2 + (\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta}) &= 0 \\ n\hat{\sigma}^2 &= (\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta}). \end{aligned} \quad (3.13.15)$$

<sup>17</sup>That our solution provides the maximum probability is easily verified in the manner demonstrated in Example ??

Solving for the constrained estimate,  $\hat{\sigma}_c^2$ , is quite simple in that it amounts to substituting  $\hat{\beta}_c$  for  $\hat{\beta}$  in Equation ???. Thus,

$$n\hat{\sigma}_c^2 = (\mathbf{Y} - \mathbf{X}\hat{\beta}_c)'(\mathbf{Y} - \mathbf{X}\hat{\beta}_c). \quad (3.13.16)$$

However, in order for it to be usable in a convenient way, we must simplify it and this is quite involved.

We'll start our simplification of  $\hat{\sigma}_c^2$  by first noting that

$$\begin{aligned} (\mathbf{Y} - \mathbf{X}\hat{\beta}_c) &= \mathbf{Y} - \mathbf{X}\hat{\beta} \\ &\quad - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \end{aligned}$$

Thus,

$$\begin{aligned} &(\mathbf{Y} - \mathbf{X}\hat{\beta}_c)'(\mathbf{Y} - \mathbf{X}\hat{\beta}_c) \\ &= (\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta}) \\ &\quad - (\mathbf{Y} - \mathbf{X}\hat{\beta})' \\ &\quad \left[ \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \right] \\ &\quad - \left[ \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \right]' \\ &\quad (\mathbf{Y} - \mathbf{X}\hat{\beta}) \\ &\quad + \left[ \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \right]' \\ &\quad \left[ \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \right]. \end{aligned}$$

While this is formidable, the second and third terms equal zero<sup>18</sup> and the last term reduces to<sup>19</sup>

$$(\boldsymbol{\theta} - \mathbf{C}\hat{\beta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\boldsymbol{\theta} - \mathbf{C}\hat{\beta}). \quad (3.13.17)$$

If we multiply Equation ??? by -1 twice, (that is, by 1), we get the equivalent and more common form,

$$(\mathbf{C}\hat{\beta} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\mathbf{C}\hat{\beta} - \boldsymbol{\theta}).$$

Thus,

$$\begin{aligned} n\hat{\sigma}_c^2 &= (\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta}) \\ &\quad + (\mathbf{C}\hat{\beta} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\mathbf{C}\hat{\beta} - \boldsymbol{\theta}) \\ &= n\hat{\sigma}^2 \\ &\quad + (\mathbf{C}\hat{\beta} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}(\mathbf{C}\hat{\beta} - \boldsymbol{\theta}) \end{aligned}$$

<sup>18</sup>To see this, all that is needed is to multiply them out and some minor cancellation. See ??? for a full derivation of this fact.

<sup>19</sup>See ??? for a full derivation of this reduction.

Now that we have MLEs for all of the parameters in the model, we can put together a likelihood ratio test.

$$\begin{aligned} \lambda &= \frac{\Pr(\mathbf{X}|\hat{\beta}_c, \hat{\sigma}_c^2)}{\Pr(\mathbf{X}|\hat{\beta}, \hat{\sigma}^2)} \\ &= \frac{(2\pi\hat{\sigma}_c^2)^{-n/2} \exp \left\{ \frac{-n(\mathbf{Y} - \mathbf{X}\hat{\beta}_c)'(\mathbf{Y} - \mathbf{X}\hat{\beta}_c)}{2(\mathbf{Y} - \mathbf{X}\hat{\beta}_c)'(\mathbf{Y} - \mathbf{X}\hat{\beta}_c)} \right\}}{(2\pi\hat{\sigma}^2)^{-n/2} \exp \left\{ \frac{-n(\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta})}{2(\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta})} \right\}} \\ &= \frac{(\hat{\sigma}_c^2)^{-n/2} \exp \frac{-n}{2}}{(\hat{\sigma}^2)^{-n/2} \exp \frac{-n}{2}} \\ &= \left( \frac{\hat{\sigma}_c^2}{\hat{\sigma}^2} \right)^{\frac{n}{2}} \end{aligned} \quad (3.13.18)$$

If Equation ??? is smaller than some constant  $k_0$ <sup>20</sup>, then the estimated variance under our hypothesis,  $\mathbf{C}\hat{\beta} = \boldsymbol{\theta}$ , would be much greater than the unconstrained estimated variance. Thus, we may infer that our hypothesis is much less accurate and less likely to be correct.

If we invert our test,

$$\left( \frac{\hat{\sigma}_c^2}{\hat{\sigma}^2} \right)^{\frac{n}{2}} \quad (3.13.19)$$

then we can reject our hypothesis when it is greater than  $k_0$ .

In order to determine what  $k_0$  is, we will use the same methods that are used to determine constants for all other likelihood ratio tests (for example, the  $z$ -test or the  $t$ -test). What we will do is set up the inequality to define the rejection region

$$\left( \frac{\hat{\sigma}_c^2}{\hat{\sigma}^2} \right)^{\frac{n}{2}} > k_0 \quad (3.13.20)$$

and use algebra to modify both sides until we can recognize the form of the left side. If the form on the left side is a standard distribution, (and in this case, since we have a ratio of variances we'll end up with an  $F$ -distribution), then we can use standard tables for this distribution to determine the value for the resulting  $k$ . Thus,

$$\begin{aligned} \left( \frac{\hat{\sigma}_c^2}{\hat{\sigma}^2} \right)^{\frac{n}{2}} &> k_0 \\ \frac{\hat{\sigma}_c^2}{\hat{\sigma}^2} &> k_1. \end{aligned}$$

<sup>20</sup>Don't worry too much about this, we'll derive the value for this constant before too long



All we have done so far is raise each side of the inequality to the  $n/2$  power. It is important to note that at this point we are no longer interested in the value for  $k_0$  but in  $k_1$ . That is why we can substitute  $k_1$  for  $k_0^{n/2}$ . Throughout this derivation we will be modifying  $k_i$  and simply replacing the old constant with a new one,  $k_{i+1}$ , since we will only be interested in the value of  $k_{i+1}$  which will eventually be determined by a standard distribution. This concept is fairly hard to grasp at first, but once you see the result, it may make more sense. Continuing with our derivation, we have,

$$\frac{n\hat{\sigma}^2 + (\mathbf{C}\hat{\beta} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1} (\mathbf{C}\hat{\beta} - \boldsymbol{\theta})}{n\hat{\sigma}^2} > k_1$$

$$1 + \frac{(\mathbf{C}\hat{\beta} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1} (\mathbf{C}\hat{\beta} - \boldsymbol{\theta})}{(\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta})} > k_1$$

$$\frac{(\mathbf{C}\hat{\beta} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1} (\mathbf{C}\hat{\beta} - \boldsymbol{\theta})}{(\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta})} > k_2$$

and

$$\frac{(\mathbf{C}\hat{\beta} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1} (\mathbf{C}\hat{\beta} - \boldsymbol{\theta})/t}{(\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta})/(n-p)} > k_3, \quad (3.13.21)$$

where  $t$  is the number of tests (or rows in  $\mathbf{C}$ ),  $n$  is the total number of observations (or rows in  $\mathbf{X}$  or  $\mathbf{Y}$ ), and  $p$  is the number of parameters we are estimating (or columns in  $\mathbf{X}$ ). We can show that Equation ?? is an  $F$ -distribution<sup>21</sup> by showing that the numerator and the denominator are both  $\sigma^2$  times chi-square variables divided by their degrees of freedom. Intuitively,  $(\mathbf{C}\hat{\beta} - \boldsymbol{\theta})$  and  $(\mathbf{Y} - \mathbf{X}\hat{\beta})$  are random variables with normal distributions in the numerator and denominator respectively. These appear as quadratic forms (that is, in the form  $\mathbf{ABA}'$ ), and thus squares of these variables. The square of any  $N(0, 1)$  random variable has a chi-square distribution. See Appendix ?? for a more complete treatment of this proof.

Since we are able to show that the left hand side of Equation ?? has an  $F$ -distribution, we can use a standard table to determine the value of  $k_3$ . That is,  $k_3 = F_{m,n-p,\alpha}$ , where  $\alpha$  is the probability we are willing to risk rejecting the hypothesis even when it is true.<sup>22</sup>

<sup>21</sup>An  $F$ -distribution is defined as the ratio of two independent chi-square variables, each divided by its degrees of freedom. That is,

$$F_{m,n} = (u/m)/(v/n),$$

where  $u \sim \chi_m^2$ ,  $v \sim \chi_n^2$  and  $u$  and  $v$  are independent.

<sup>22</sup>Just as a gentle reminder, the value,  $F_{m,n-p,\alpha}$  and the

### Example 3.13.5.3

Given the Leghorn data in Table ?? we can try to fit the model,

$$y = \beta_0 + \beta_1 x,$$

to the data. Thus, setting up our  $\mathbf{Y}$  and  $\mathbf{X}$  matrices we have,

$$\mathbf{Y} = \begin{bmatrix} 87.1 \\ 93.1 \\ 89.8 \\ 91.4 \\ 99.5 \\ 92.1 \\ 95.5 \\ 99.3 \\ 93.4 \\ 94.4 \end{bmatrix} \quad \text{and} \quad \mathbf{X} = \begin{bmatrix} 1 & 4.6 \\ 1 & 5.1 \\ 1 & 4.8 \\ 1 & 4.4 \\ 1 & 5.9 \\ 1 & 4.7 \\ 1 & 5.1 \\ 1 & 5.2 \\ 1 & 4.9 \\ 1 & 5.1 \end{bmatrix}$$

Using the Octave program listed in Appendix ??, which simply implements the test developed in Section ??, we can test different hypotheses.

To test whether or not body weight is an indicator of food consumption, that is, whether or not  $\beta_1 = 0$ , we use the contrast matrix:

$$\mathbf{C} = [0, 1].$$

Our constraint for estimating the parameters is thus,

$$\mathbf{C}\boldsymbol{\beta} = 0$$

$$\begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix} = 0$$

$$\beta_1 = 0$$

and the output from our program is:

```
octave:2> general_linear (x, y, [0, 1]);
beta =

55.2633
7.6901

f_test = 16.232
p_value = 0.0037939
t_test = 4.0289
```

Thus, since the p-value is much smaller than 0.05, we can conclude that  $\beta_1 \neq 0$  and that body weight does indeed give an indication of food consumption. ||

### Example 3.13.5.4 (ANOVA)

0 ppm	62.5 ppm	250 ppm	1000 ppm
55	47	49	36
47	51	44	41
46	40	46	
53	44	51	
		47	

Table 3.13.2: Study effect of different levels of PCB on body weights (in grams) of mice. *Source:* Plagiarized from Roger L. Berger, [?]

Given the data in Table ??, a modern ANOVA model<sup>23</sup>

$$y = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_4 x_4,$$

is easily computed using Equation ?. We can set up our matrices as follows:

$$\mathbf{Y} = \begin{bmatrix} 55 \\ 47 \\ 46 \\ 53 \\ 47 \\ 51 \\ 40 \\ 44 \\ 49 \\ 44 \\ 46 \\ 51 \\ 48 \\ 36 \\ 41 \end{bmatrix} \quad \text{and} \quad \mathbf{X} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

To test the hypothesis that there is no difference in the treatments (that is, all of the parameters are equal), we create the contrast matrix:

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{bmatrix}.$$

ratio in Equation ?? represent points on an  $x$ -axis. The value,  $F_{m,n-p,\alpha}$ , represents a cut-off point, and anything larger, and thus further away from the mean, is determined to not come from the same distribution.

<sup>23</sup>This use of the word *modern* is perhaps wishful thinking as it is the author's opinion that this model should be considered thus. In practice, most people, for historical reasons, use alternative models for ANOVA. See Appendix ??, for a full discussion on this topic.

The constraint on our parameter estimation is thus,

$$\mathbf{C}\boldsymbol{\beta} = 0$$

$$\begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$\begin{bmatrix} \beta_1 - \beta_4 \\ \beta_2 - \beta_4 \\ \beta_3 - \beta_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

Using our Octave program, we get:

```
> [x, y] = load_data("mice.dat");
> c = [1, 0, 0, -1;
> 0, 1, 0, -1; 0, 0, 1, -1]
c =
```

```
1 0 0 -1
0 1 0 -1
0 0 1 -1
```

```
> general_linear(x, y, c);
beta =
```

```
50.250
45.500
47.600
38.500
```

```
f_test = 4.3057
p_value = 0.030746
```

Thus, since our p-value is less than 0.05, we can reject our hypothesis that there is no difference in the treatment means. ||

### Example 3.13.5.5 (Randomized Block Design)

Sometimes there are extra sources of variation in your data that you can clearly define. For example, in a laboratory, 5 different people may have run the same experiment and recorded their results. We could treat all of the experiments equally, ignoring the fact that we know that 5 different people were involved and just use standard ANOVA for hypothesis testing, or we could try to take into account the fact that some people may be more precise than others when making measurements and that some of the variation in the data may be due to this. The goal of this example is to demonstrate how this second option can work and how we can effectively make use of the extra information at hand.

Before we begin, however, we will note that this method only makes sense in specific situations and that, with a little more data, we can use much more powerful methods for analysis. In this example we are assuming that it would be unreasonable to have each person run the experiment more than once. If this were not the case, we would apply the methods explained in Example ??.

Consider the data set in Table ??, where we have a single experiment to compare three different methods for drying leaves after they have been rinsed performed by five different people. The term *Block* refers

Treatment	Person Applying Treatment (Block)				
	1	2	3	4	5
Control	950	887	897	850	975
Blotted	857	1,189	918	968	909
Air Current	917	1,072	975	930	954

Table 3.13.3: Study of drying methods (measured in seconds till try) *Source:* Plagiarized from Steel, Torrie and Dickey, [?]. Data from Tucker et al.

to how randomization was used to set up the experimental design. Here we are considering each person to represent a block and that within each block, the order in which the different drying methods is applied is randomized. That is to say, Person (Block) #1 may have used the blotting method first, then the control method and then the air current method. Person (Block) #2 may have started with the control method, then used the air current method and then the blotting method. This randomization within blocks is important to make sure that bias is not introduced to the data by where a treatment is applied. It could be that if one treatment was always done last, then it could have a bias due to the student wanting to get done quickly and go home for the evening.

For the data in Table ??, the model is:

$$y = \beta_1 b_1 + \beta_2 b_2 + \beta_3 b_3 + \beta_4 b_4 + \beta_5 b_5 + \beta_6 t_1 + \beta_7 t_2 + \beta_8 t_3,$$

where  $b_1, b_2, b_3, b_4,$  and  $b_5,$  are indicator variables (dummy variables) for which block the sample came from and  $t_1, t_2$  and  $t_3$  are indicator variables for which treatment was applied to the sample. Thus, it would

make sense for  $\mathbf{Y}$  and  $\mathbf{X}$  to be:

$$\mathbf{Y} = \begin{bmatrix} 950 \\ 857 \\ 917 \\ 887 \\ 1,189 \\ 1,072 \\ 897 \\ 918 \\ 975 \\ 850 \\ 968 \\ 930 \\ 975 \\ 909 \\ 954 \end{bmatrix} \quad \mathbf{X} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix},$$

and while  $\mathbf{Y}$  is fine,  $\mathbf{X}$  is a singular<sup>24</sup> matrix and we will not be able to invert  $\mathbf{X}'\mathbf{X}$ . A common solution to this problem is to use tricks that are often applied to traditional ANOVA models (ANOVA models that contain a general location parameter or overall mean and the remaining parameters are deviations from that mean) to reparameterize  $\beta$ .

<sup>24</sup>This can be seen by first adding together the first five columns in  $\mathbf{X}$ , which will give you a vector of 1s. Adding the last three columns of  $\mathbf{X}$  together also gives you a vector of 1s. Thus, adding the first five columns and subtracting the last three columns will result in a vector of 0s, satisfying the definition of a singular matrix.

**THE MAIN IDEA - Reparameterization:**

When the design matrix,  $\mathbf{X}$ , is singular, than you will need to reformulate the parameter vector,  $\beta$ , in order to make  $\mathbf{X}$  nonsingular. The reparameterization should be such that you can reduce the number of columns in  $\mathbf{X}$  to a reasonable number. It is important, however, that the new  $\mathbf{X}$  be of full rank in order to prevent overestimation of the error. That is, by estimating a less than full rank model, the residual variation is contaminated by the variations contributed by the unestimated parameters. Full rank means, in a rough sense, that the matrix has as many columns as is possible while remaining non-singular.

It should also be noted that the need to reparameterize should be a signal that you are making potentially unreasonable assumptions about the data. In this example, we only have one measurement per Person/Treatment combination and thus, can not estimate the error within these combinations. This lack of data is what forces us to reparameterize and assume that this error is insignificant.

In order to reparameterize  $\beta$ , we will consider there to be a common mean for all of the data,  $\mu$  and that the effects of different blocks and treatments represent deviations from this mean. Thus, each data point  $y_{i,j} = \mu + \Delta b_{i,j} + \Delta t_{i,j}$ . It is also important to note that due to Equation ?? (where we show that our estimates in  $\beta$  are unbiased), the sum of the block variations and the sum of the treatment variations are equal to zero and that any particular block variation or treatment variation can be derived from the others. That is to say,

$$\sum_{i=1}^5 \Delta b_i = 0, \quad (3.13.22)$$

and thus,

$$\Delta b_5 = -\Delta b_1 - \Delta b_2 - \Delta b_3 - \Delta b_4,$$

and

$$\begin{aligned} \sum_{i=1}^3 \Delta t_i &= 0 & (3.13.23) \\ \Delta t_1 + \Delta t_2 + \Delta t_3 &= 0 \\ \Delta t_3 &= -\Delta t_1 - \Delta t_2. \end{aligned}$$

Since we can derive both  $\Delta b_5$  and  $\Delta t_3$  from the other deviations, we do not need to estimate values for them and as such, we can omit them from  $\beta$ . By adding  $\mu$  and removing  $\Delta b_5$  and  $\Delta t_3$ ,  $\beta$  has one fewer row

than before and as a result, we can remove a column from  $\mathbf{X}$  and it will become non-singular. Thus,

$$\beta = \begin{bmatrix} \mu \\ \Delta b_1 \\ \Delta b_2 \\ \Delta b_3 \\ \Delta b_4 \\ \Delta t_1 \\ \Delta t_2 \\ \Delta t_3 \end{bmatrix}.$$

and design matrix then becomes

$$\mathbf{X} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & -1 & -1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & -1 & -1 \\ 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & -1 & -1 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & -1 & -1 \\ 1 & -1 & -1 & -1 & -1 & 1 & 0 \\ 1 & -1 & -1 & -1 & -1 & 0 & 1 \\ 1 & -1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix}.$$

To test the hypothesis that there is no difference in the treatments we want to know if the treatment deviations are all zero. We can create a contrast matrix to directly test to see if  $\Delta t_1$  and  $\Delta t_2$  are zero and if they are, by Equation ??, we can infer that  $\Delta t_3$  is also zero. Thus,

$$\mathbf{C} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Using Octave we get:

```
> general_linear (x, y, c);
beta =
    949.867
   -41.867
    99.467
   -19.867
   -33.867
   -38.067
    18.333
f_test = 0.82474
p_value = 0.47244
```

and since the p-value is greater than 0.05, we will conclude that there is no difference in the three treatments. It is worth noting that if we had ignored the blocks and just used regular ANOVA, than the f-statistic would have been 0.71 and the p-value would have been 0.5128, demonstrating that by not ignoring the blocks, we have increased the precision of our test, even though we would have come to the same conclusion.

If we wanted to test the hypothesis that there was no difference in the different blocks (in this case this would mean that each person applied the treatment in more or less the same way), than we would have to test to see if the block deviations were all zero. Using the same logic used to create the contrast matrix for the treatment hypothesis we have:

$$C = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}.$$

Using Octave to calculate an F-statistic and a p-value, we get:

```
general_linear (x, y, c);
beta =

    949.867
   -41.867
    99.467
   -19.867
   -33.867
   -38.067
    18.333
```

```
f_test = 1.5022
p_value = 0.28882
```

and once again, we would fail to reject our hypothesis and conclude that there is no significant difference between the blocks.

At this point it is worth noting that if we had more than one measurement per person/drying method combination, then we would not have had to do the reparameterization. With only one measurement per combination, we are forced to assume that there is no interaction between the person and the drying method. With more measurements, we do not have to make this assumption and we would have treated the data as we would a 5x3 factorial experiment. How this is done is shown in Examples ??, ?? and ??.

### THE MAIN IDEA - Block Designs:

Any design that includes blocking also includes assumptions that some of the potential interactions do not exist. In Example ??, it could be that some of the students had a lot of experience with one of the drying methods. By using blocking, we are simply assuming that this can not be the case. Thus, if at all possible, you must try to obtain more than one measurement per combination of factors.

### Example 3.13.5.6 (2x2 Factorial)

Maternal Diet	Age			
	Adolescent		Mature	
0%	5	4	6	7
	3	4	5	8
	2		4	
35%	18	19	6	9
	14	12	5	9
	15		3	

Table 3.13.4: Study of mouse learning times, testing both how maternal diet (calories derived from ethanol) and age might effect how many times a mouse repeats a test before passing. *Source:* Plagiarized from M. Plonsky. *Analysis of Variance - Two Way* <http://www.uwsp.edu/psych/stat/13/anova-2w.htm>

Maternal Diet	Age	
	Adolescent	Mature
0%	3.6	6
35%	15.6	6.4

Table 3.13.5: Average value for each factor combination.

Sometimes you are interested in testing more than one variable at a time. For example, you might want to do this to determine if there are any interactions between two drugs at different dosages. Such an interaction might prevent the two drugs from being as effective together as they would be taken alone. A doctor might find this information helpful when deciding how to make prescriptions. In this case, you would need to use some type of factorial method to test your hypothesis. In this case, the the different drugs would be called Factors and the different dosages of each drug would be called Levels. If

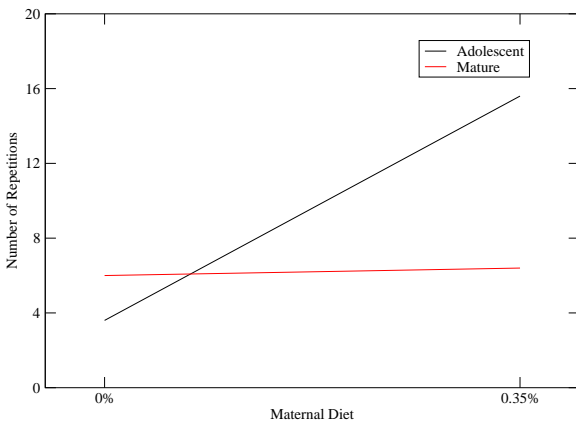


Figure 3.13.3: A plot of the averages for the combinations of factors. That the lines are not parallel indicates that there is an interaction between the two factors.

we are interested in only two different dosages per drug, then we would use what is called a 2x2 factorial method (2x2 indicating that there is a total of 4 different drug/dosage combinations). In general, factorial method can be used to determine if there is any interactions between the two factors and if there are none, it can then determine the significance of the effects of the different levels of the factors.

In a completely randomized experiment with multiple factors, there is no attempt to impose any type of blocking structure. That is to say, the randomization occurs at the top most level of organization (i.e. we randomly select a level from one factor and apply it with a randomly selected level from another factor). In a blocking design, we first designate the blocks and only randomize within the blocks.

The general model used for a 2x2 factorial method

$$y_i = \beta_1(a_1b_1)_i + \beta_2(a_1b_2)_i + \beta_3(a_2b_1)_i + \beta_4(a_2b_2)_i,$$

where  $(a_xb_y)_i$  are indicator variables that show that the measured value came from a particular combination of factors.

Given the data in Table ??, we can apply the 2x2 factorial method to determine if there is any interaction between a mouse's age and how much prenatal ethanol its mother consumed when determining how quickly it can learn something new. We might speculate that older mice are going to be slow learners regardless of their prenatal environment, where as young mice might be heavily influenced this factor. We can further validate this hunch by creating what is called an *interaction plot*, Figure ??, by plotting the average responses for the adolescent and mature

mice given the two different conditions (data from Table ??). In this case we see the two lines are not parallel and this indicates that the age of the mouse potentially changes the size of the effect that the pre-natal environment has on its ability to learn.<sup>25</sup>

Since the model we will be using is a linear model, we can apply our general method for hypothesis testing. Thus, the matrices for  $\mathbf{Y}$  and  $\mathbf{X}$  are:

$$\mathbf{Y} = \begin{bmatrix} 5 \\ 4 \\ 3 \\ 4 \\ 2 \\ 6 \\ 7 \\ 5 \\ 8 \\ 4 \\ 18 \\ 19 \\ 14 \\ 12 \\ 15 \\ 6 \\ 9 \\ 5 \\ 9 \\ 3 \end{bmatrix} \quad \mathbf{X} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

To test if there is interaction effects between the factors, and Figure ?? indicates that this is probably the case, we want to test if the slopes in the plot are equal. That is,  $(\beta_1 - \beta_2) - (\beta_3 - \beta_4) = \beta_1 - \beta_2 - \beta_3 + \beta_4 = 0$  and thus, we define the contrast matrix,

$$\mathbf{c} = [ 1 \quad -1 \quad -1 \quad 1 ],$$

and our Octave program gives us the results:

```
> general_linear (x, y, c);
beta =

    3.6000
   15.6000
    6.0000
    6.4000

f_test = 35.598
p_value = 1.9738e-05
t_test = 5.9664
```

<sup>25</sup>A large number of different interaction plots, as well as their potential interpretations is given in Appendix ??.




Figure 3.13.4: A plot of the averages for the combinations of factors. To test for interactions, we would need to determine that the line segments between each level of factor A are parallel.

The tiny p-value confirms the intuition gained from Figure ??.

### Example 3.13.5.7 (NxM Factorial)

If there are more than two levels per factor, then the only thing that changes is the number of columns in the design matrix, to account for the larger number of possible combinations between factors, and the number of rows in the contrast matrix.

For example, if you were to do a 3x2 factorial,  $\mathbf{X}$  would have 6 columns and in order to test for interactions,  $\mathbf{C}$  would require 2 rows. The need for the extra row in  $\mathbf{C}$  is illustrated in Figure ??, where the two groups of line segments would need to be tested to see if their slopes are equal.

### Example 3.13.5.8 (2x2x2 Factorial)

By adding an additional factor, you add an extra dimension to types of interactions that can take place. Instead of a two dimensional table to describe the combinations of different levels of different factors, you now need a three dimensional cube. Additional factors require additional dimensions. Fortunately all of these extra dimensions can be represented by multiple two dimensional tables.

Consider the data in Table ?. If we let Body Fat to be factor A, Gender to be factor B and Smoking History to be factor C, we can represent this data

with the model:

$$y_i = \beta_1(a_1b_1c_1)_i + \beta_2(a_1b_1c_2)_i + \beta_3(a_1b_2c_1)_i + \beta_4(a_1b_2c_2)_i + \beta_5(a_2b_1c_1)_i + \beta_6(a_2b_1c_2)_i + \beta_7(a_2b_2c_1)_i + \beta_8(a_2b_2c_2)_i;$$

where  $(a_ib_jc_k)_i$  is 1 if the data came from that combination of factor levels.

		Smoking History	
		Light	Heavy
Low Body Fat	Male	24.1	17.6
		29.2	18.8
	Female	24.6	23.2
		20.0	14.8

		Smoking History	
		Light	Heavy
High Body Fat	Male	14.6	14.9
		15.3	20.4
	Female	12.3	12.8
		16.1	10.1

Table 3.13.6: Study of how Body Fat, Smoking and Gender effect how long it takes to reach fatigue on an exercise bike. Source: Plagiarized from Neter et al., [?]

		Smoking History	
		Light	Heavy
Low Body Fat	Male	25.967	19.867
	Female	19.833	12.133

		Smoking History	
		Light	Heavy
High Body Fat	Male	14.067	16.033
	Female	12.067	10.2

Table 3.13.7: Average value for each Factor/Level combination.

Now that we have the model, the question becomes, how do we test hypotheses about it. Can we use it to determine if Body Fat has an effect of its own on how quickly people get fatigued or is it confounded with the other factors? What about smoking? Do differences in how quickly someone is fatigued depend

on all three factors working simultaneously? With carefully formulated contrast matrices we can answer all of these questions.

Before we start, however, it is important to understand that you need to begin by looking for the highest order interactions (in this case, the potential three way interaction between the three factors) before testing hypotheses about lower order interactions, which, in turn, need to be investigated before you look at whether or not any particular factor has a main effect (operates on its own). This is because if there are higher order interactions involving the specific factor you are interested in testing for main effects, then the data set is not usable for the kind of hypothesis you have in mind. Despite this, we begin by formulating the contrast matrices to test for main effects. The reason we do this, however, is that once we have the matrices needed to test for main effects, we can use them to derive all of the other contrast matrices required test for any possible higher order interactions. Thus, even though we start by creating the contrasts to test for main effects, we do not use them individually until the very end of the analysis if we use them at all.

Consider the data from Table ?? as it is plotted in Figure ?. If we assumed that there were no higher order interactions, we could test to determine if Body Fat has a main effect by testing to see if the slopes of the lines between Low and High were zero. We can do this by the usual way by calculating the slopes and then determining if they are significantly different from zero or not. Thus, the contrast matrix for the hypothesis that there is no main effect (the slopes are zero) is:

$$\mathbf{C} = [ -1 \quad -1 \quad -1 \quad -1 \quad 1 \quad 1 \quad 1 \quad 1 ] .$$

Similarly, to test for whether or not Gender has a main effect, we would want to test for whether or not the slopes in Figure ?? were zero or not. The hypothesis that there is no main effect is encoded into the contrast matrix:

$$\mathbf{C} = [ -1 \quad -1 \quad 1 \quad 1 \quad -1 \quad -1 \quad 1 \quad 1 ] .$$

Figure ?? implies that in order to test for a main effect from Smoking, we must define the contrast matrix as follows:

$$\mathbf{C} = [ -1 \quad 1 \quad -1 \quad 1 \quad -1 \quad 1 \quad -1 \quad 1 ] .$$

Now, all that we need to do to test hypothesis about any higher order interactions is multiply, column by column, the contrast matrices associated

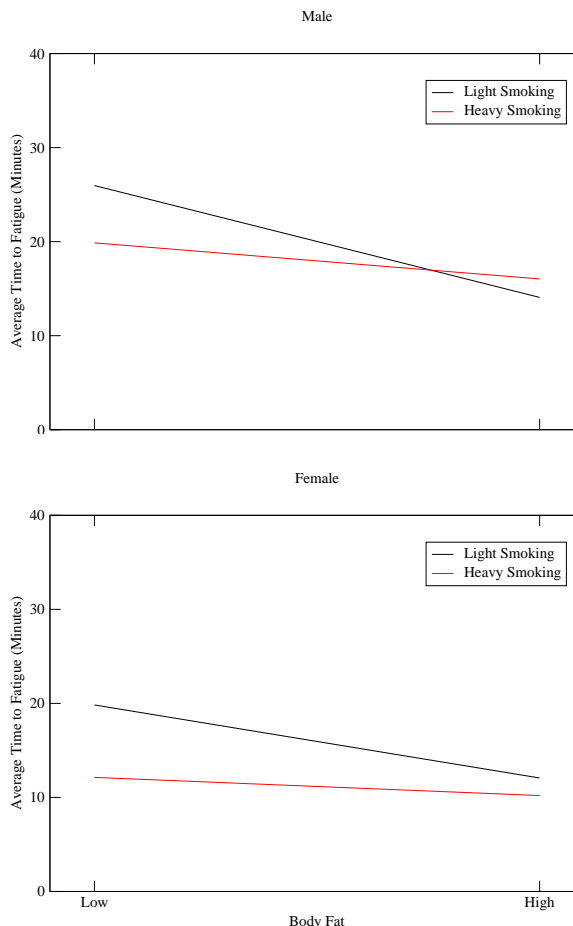


Figure 3.13.5: Interaction plots for Body Fat and Smoking.



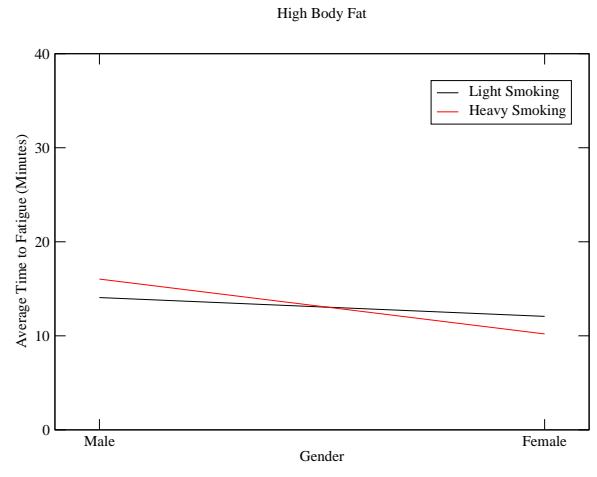





Figure 3.13.6: Interaction plots for Gender and Smoking.


Figure 3.13.7: Interaction plots for Smoking and Gender.

with each factor in the interaction. For example, to test for an interaction between Gender and Smoking, we create the new contrast matrix

$$\begin{aligned}
 & \begin{bmatrix} -1 & -1 & 1 & 1 & -1 & -1 & 1 & 1 \end{bmatrix} \\
 \mathbf{C} = & \begin{bmatrix} -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 \end{bmatrix} \\
 & \begin{bmatrix} 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \end{bmatrix}.
 \end{aligned}$$

This contrast matrix is equivalent to testing whether the slopes of the lines within each graph in Figure ?? are equal.


1.



To test a hypothesis about a three way interaction between Body Fat, Gender and Smoking, we simply multiply the columns of all three main effect matrices together. This contrast matrix is

$$\begin{aligned}
 & \begin{bmatrix} -1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 \end{bmatrix} \\
 \mathbf{C} = & \begin{bmatrix} -1 & -1 & 1 & 1 & -1 & -1 & 1 & 1 \end{bmatrix} \\
 & \begin{bmatrix} -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 \end{bmatrix} \\
 & \begin{bmatrix} -1 & 1 & 1 & -1 & 1 & -1 & -1 & 1 \end{bmatrix}.
 \end{aligned}$$

2.



Intuitively, one can also imagine that this contrast matrix tests to determine if the slopes between the graphs in Figure ?? are equal (see Figure ??).

As before, to test our hypotheses we create  $\mathbf{Y}$  and

Figure 3.13.8: The two graphs in Figure ?? can be thought of as forming the front and back sides of a cube. To test for a two way interaction between Gender and Smoking, you simply need to verify that the two lines *within* the front and back sides are parallel (1). To determine if there is a three way interaction between Gender, Smoking and Body Fat, you have to verify that the lines drawn *between* the front and the back are parallel (2). The lines traveling between the front and the back effectively traverse all three dimensions of the space defined by the three factors. The lines within the front and back sides only make use of two dimensions.

X

$$\mathbf{Y} = \begin{bmatrix} 24.1 \\ 29.1 \\ 24.6 \\ 17.6 \\ 18.8 \\ 23.2 \\ 20.0 \\ 21.9 \\ 17.6 \\ 14.8 \\ 10.3 \\ 11.3 \\ 14.6 \\ 15.3 \\ 12.3 \\ 14.9 \\ 20.4 \\ 12.8 \\ 16.1 \\ 9.3 \\ 10.8 \\ 10.1 \\ 14.4 \\ 6.1 \end{bmatrix} \quad \mathbf{X} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Using our contrast matrix for testing for three way interactions between Body Fat, Gender and Smoking,

$$\mathbf{C} = [ -1 \ 1 \ 1 \ -1 \ 1 \ -1 \ -1 \ 1 ],$$

we can use our Octave program and the output is:

```

> general_linear(x, y, c);
beta =

25.967
19.867
19.833
12.133
14.067
16.033
12.067
10.200

f_test = 0.20036
p_value = 0.66043
t_test = 0.44761

```

Since the p-value is much larger than 0.05, we will fail to reject the hypothesis that there is no three way interaction between Body Fat, Gender and Smoking.

That is to say, there is no three way interaction between the three factors in the study. Thus, we are free to test hypotheses about any possible two way interactions that might be present.

Moving on, we now test for interactions between Gender and Smoking. As described above, we define the contrast matrix:

$$\mathbf{C} = [ 1 \ -1 \ -1 \ 1 \ 1 \ -1 \ -1 \ 1 ],$$

and octave gives us:

```

> general_linear(x, y, c);
beta =

25.967
19.867
19.833
12.133
14.067
16.033
12.067
10.200

f_test = 1.1859
p_value = 0.29230
t_test = 1.0890

```

leaving us to conclude that, since the p-value is greater than 0.05, there are no interactions between Gender and Smoking.

To test if there are interactions between Body Fat and Smoking,

$$\mathbf{C} = [ 1 \ -1 \ 1 \ -1 \ -1 \ 1 \ -1 \ 1 ],$$

and octave gives us:

```

> general_linear(x, y, c);
beta =

25.967
19.867
19.833
12.133
14.067
16.033
12.067
10.200

f_test = 7.7612
p_value = 0.013221
t_test = 2.7859

```

so that we would reject the hypothesis that there is no interaction between Body Fat and Smoking.

To test for interactions between Gender and Body Fat,

$$C = \begin{bmatrix} 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \end{bmatrix},$$

and octave gives us:

```
> general_linear(x, y, c);
beta =
```

```
25.967
19.867
19.833
12.133
14.067
16.033
12.067
10.200
```

```
f_test = 1.4622
p_value = 0.24414
t_test = 1.2092
```

and we would fail to reject the hypothesis that there is no interaction between Gender and Body Fat.

Finally, since we have determined that Gender is the only factor that is not in a two interaction, we can determine if has a main effect. To test if the hypothesis is that there is no main effect, the contrast matrix is, as defined above,

$$C = \begin{bmatrix} -1 & -1 & 1 & 1 & -1 & -1 & 1 & 1 \end{bmatrix},$$

and octave gives us:

```
> general_linear(x, y, c);
beta =
```

```
25.967
19.867
19.833
12.133
14.067
16.033
12.067
10.200
```

```
f_test = 18.915
p_value = 0.00049705
t_test = 4.3492
```

The small p-value tells us to reject our hypothesis and conclude that there is a main effect for Gender.

||

### Example 3.13.5.9 (Blocked 2x3 Factorial)

It is possible to combine blocking with the factorial method for analyzing data. Using blocking implies that you are assuming that some of the potential interactions do not exist while leaving room for the possibility with other interactions.

Consider the data in Table ???. Here we have two factors, Major and Grade, that have been separated into different blocks defined by Teacher (the measured value is Score). Notice that just like in Example ??, we only have one measurement for each cell in each block. This prevents us from using the full models from Examples ?? and ?? that allowed us to test for all possible interactions. Instead, we are forced to assume that there is no three way interaction between Teacher, Major and Grade. However, since we have multiple measurements for each Major/Grade combination, we can test for two way interactions between these two factors.

Teacher (Block)	Grade			
	Major	Jr.	Sr.	Gr
#1	CS	80	80	80
	MA	85	90	96
#2	CS	75	80	84
	MA	80	88	97
#3	CS	75	80	85
	MA	75	80	100

Table 3.13.8: Data for a 2x3 Factorial with Blocking. Teacher is used for blocking and Major and Grade are the two factors. The measured value is Score. *Source:* Plagiarized from Tanya Balan.

Given the restrictions of the data, the model is thus,

$$y = \beta_1 t_1 + \beta_2 t_2 + \beta_3 t_3 + \beta_4 m_1 + \beta_5 m_1 + \beta_6 g_1 + \beta_7 g_2 + \beta_8 g_3 + \beta_9(m_1 g_1) + \beta_9(m_1 g_2) + \beta_9(m_2 g_1) + \beta_9(m_2 g_2),$$

where  $t_i$ ,  $m_j$  and  $g_k$  are the indicator variables for Teacher, Major and Grade, respectively, and  $(m_j g_k)$  are the indicator variables for the different types of two way interactions.

As with our previous experience with blocking in Example ??, this equation will lead to a singular design matrix and thus we must reparameterize in much the same way we did before. The only significant difference is how we rewrite the interaction terms. Here, we will use the methods we learned in Example ??, where we built the interaction contrasts from the main effect contrasts. Instead of multiplying each column in contrast matrices, we will multiply the rows that determine which specific combination of Major and Grade a particular score is associated with to create the columns that specify the interaction. Thus, the design matrix is:

$$\mathbf{X} = \begin{bmatrix} 1 & 1 & 0 & 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 1 & -1 & -1 & -1 & -1 \\ 1 & 1 & 0 & -1 & 1 & 0 & -1 & 0 \\ 1 & 1 & 0 & -1 & 0 & 1 & 0 & -1 \\ 1 & 1 & 0 & -1 & -1 & -1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & 0 & 1 & -1 & 1 & 0 & -1 & 0 \\ 1 & 0 & 1 & -1 & 0 & 1 & 0 & -1 \\ 1 & 0 & 1 & -1 & -1 & -1 & 1 & 1 \\ 1 & -1 & -1 & 1 & 1 & 0 & 1 & 0 \\ 1 & -1 & -1 & 1 & 0 & 1 & 0 & 1 \\ 1 & -1 & -1 & 1 & -1 & -1 & -1 & -1 \\ 1 & -1 & -1 & -1 & 1 & 0 & -1 & 0 \\ 1 & -1 & -1 & -1 & 0 & 1 & 0 & -1 \\ 1 & -1 & -1 & -1 & -1 & -1 & 1 & 1 \end{bmatrix},$$

where Column 1 is the overall mean,  $\mu$ , the Columns 2-3 characterize which Teacher, or block, the data comes from, Column 4 specifies which Major the data comes from, Columns 5-6 determines which Grade the data comes from and Columns 7-8, the product of Column 4 and Columns 5-6, specifies the (Major  $\times$  Grade) interaction.

As with the factorial method, we need to evaluate the presence of interaction before we determine the effects of the individual factors. To do this, we make the null hypothesis that there is no interaction. This is equivalent to assuming the parameters associated with Columns 7 and 8 are both equal to zero. The

contrast matrix is thus,

$$\mathbf{C} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Using our Octave program, we test our hypothesis:

```
> general_linear (x, y, c);
beta =

84.16667
1.00000
-0.16667
-4.27778
-5.83333
-0.33333
2.61111
0.44444
```

```
f_test = 5.3515
p_value = 0.026292
```

and, due to the small p-value, reject it.

Since we have concluded that there is indeed interactions between Major and Grade, we can end our inquiry here. ||

### Example 3.13.5.10 (Split-plot Design)


Split-plot designs are simply designs that incorporate multiple levels of blocking, and thus, only attempt to model a subset of the potential sources of variation. The more “split” you see in the name, like split-split-plot, of the design, the more levels of blocking there are in it. In this case, where we are just demonstrating split-plot, we let one factor determining the blocking at the top level, these blocks are then divided into two sub-blocks, or plots, with the levels of a second factor applied to each sub-block. These sub-blocks are then divided into sub-sub-blocks and the levels of a third factor are randomly distributed among these units. This is illustrated in Figure ??.

Conceptually, however, they do offer anything we have not seen before.


Consider the data in Table ??. The table alone is enough to convince us that any hope we might have to model potential three way interactions is a lost cause. The fact that there is only one measurement per cell rules this out by not giving us enough degrees of freedom to model both a three way interaction and estimate the error. We are forced to simply assume that this type of interaction can not exist.

On the other hand, we can model a two way interaction between Field and Variety because we have

1.



2.



3.




Figure 3.13.9: Assume there are three Factors, A, B and C, each with 2 Levels. First, (1) we divide the levels of A between two blocks, randomly selecting which level goes with which block. Second, (2) we divide each block into two sub-blocks and divide the levels of B between them by randomly selecting which level goes with which sub-block. Lastly, (3) these sub-blocks are divided and the levels of C are randomly distributed among these subunits.

Field #1

Variety	Row Spacing				
	18	24	30	36	42
OM	33.6	31.1	33.0	28.4	31.4
B	28.0	23.7	23.5	25.0	25.7

Field #2

Variety	Row Spacing				
	18	24	30	36	42
OM	37.1	34.5	29.5	29.9	28.3
B	25.5	26.2	26.8	25.3	23.2

Field #3

Variety	Row Spacing				
	18	24	30	36	42
OM	34.1	30.5	29.2	31.6	28.9
B	28.3	27.0	24.9	25.6	23.4

Field #4

Variety	Row Spacing				
	18	24	30	36	42
OM	34.6	30.5	29.2	31.6	28.9
B	29.4	25.8	23.3	26.4	25.6

Field #5

Variety	Row Spacing				
	18	24	30	36	42
OM	35.4	30.7	30.7	28.1	18.5
B	27.3	26.8	21.4	24.6	24.5

Field #6

Variety	Row Spacing				
	18	24	30	36	42
OM	36.1	30.3	27.9	26.9	33.4
B	28.3	23.8	22.0	24.5	22.9

Table 3.13.9: Yield, in bushels per acre. There were six different fields (Blocks) and these blocks were divided into two sub-blocks, each sub-block receiving one of two different varieties of soybean. Within each sub-block, five different row spacings were applied to the rows of soybeans. Source: Plagiarized from Steel, Torrie and Dickey, [?]. Data from J. W. Lambert, University of Minnesota

multiple measurements per level of Variety in each Field. We can also model a two way interaction between Variety and Spacing because, across the different Fields, we can come up with several measurements for each Variety/Spacing combination. Thus, the model is:

$$y = \beta_i \text{Field} + \beta_j \text{Variety} + \beta_k \text{Spacing} + \beta_l (\text{Field} \times \text{Variety}) + \beta_m (\text{Variety} \times \text{Spacing}).$$

The more common form of this model is:

$$y = \beta_i \text{Field} + \beta_j \text{Variety} + \beta_l (\text{Field} \times \text{Variety}) + \beta_k \text{Spacing} + \beta_m (\text{Variety} \times \text{Spacing}).$$

Since we do not have enough degrees of freedom to estimate all of the parameters in a full model (i.e. we are leaving out the three way interaction term) then we will have to reparameterize the design matrix in the same way we did in Example ???. This gives us a design matrix with sixty rows and twenty columns (thank goodness someone has already gone ahead and written a computer program to generate these for us, [?]). Since this is quite large, I will leave it to your imagination, however, a small bit of it can be found in Appendix ???.

Analysis of the data follows that of any model that contains interaction terms. Start by analyzing those to determine if there is interaction. If not, then test for main effects. If there is, then, you have done your best.

The contrast matrix for testing for interaction between Variety and Spacing can be found in Appendix ??? and the result of our test is:

```
> general_linear (x, y, c);
beta =
```

```
28.111667
 0.228333
 0.518333
 0.238333
 0.828333
-1.311667
 2.821667
 0.338333
 0.408333
-0.311667
 0.018333
-0.941667
 3.363333
 0.480000
-1.203333
```

```
-0.728333
 0.853333
 0.220000
 0.436667
-0.671667
```

```
f_test = 1.3051
p_value = 0.28457
```

and the large p-value lets us conclude that there is no interaction between Variety and Spacing.

Since spacing is not tied up in another interaction term, like Variety, we can go ahead and test for whether or not it has a main effect. This contrast can be found in Appendix ??? and octave tells us:

```
> general_linear (x, y, c);
beta =
```

```
28.111667
 0.228333
 0.518333
 0.238333
 0.828333
-1.311667
 2.821667
 0.338333
 0.408333
-0.311667
 0.018333
-0.941667
 3.363333
 0.480000
-1.203333
-0.728333
 0.853333
 0.220000
 0.436667
-0.671667
```

```
f_test = 10.567
p_value = 6.1454e-06
```

Such a small p-value (far less than 0.05) allows us to conclude that Spacing does indeed have a main effect on the yield.

We will now test to see if the interaction term Field×Variety is significant. The contrast matrix for this can be found in Appendix ???. The results of our test are:

```
> general_linear (x, y, c);
beta =
```

28.111667  
 0.228333  
 0.518333  
 0.238333  
 0.828333  
 -1.311667  
 2.821667  
 0.338333  
 0.408333  
 -0.311667  
 0.018333  
 -0.941667  
 3.363333  
 0.480000  
 -1.203333  
 -0.728333  
 0.853333  
 0.220000  
 0.436667  
 -0.671667

f\_test = 0.61686  
 p\_value = 0.68760

and once again, since the p-value is so large, we will fail to reject the hypothesis that the interaction term is insignificant.

Now that we have established that both Field and Variety are free of any interaction, we can test to see if they have any main effects. As we might expect Field ends up not having a main effect, but Variety does.

It is worth noting that often times statisticians will modify the formula used to calculate the F-statistic when they are testing for main effects for factors whose levels are not randomly applied within the sub-blocks. Instead of using the Mean Square Error (the estimation for the overall error in the model) in the denominator, they will use the mean square of Block×Plot interaction term. Doing so has the potential to increase the power of this test. ||

### Example 3.13.5.11 (ANCOVA)

Analysis of covariance (ANCOVA) is very similar to ANOVA in that the models contain indicator variables for the various treatments. The differences comes from the fact that ANCOVA models also contain variables that represent continuous data, like weight or height. These continuous variables are called *covariates* because, even though they are not

controlled by the experimenter, they are expected to have an influence on the response to the treatments. You can imagine how heavy and light people might have different response to different “treatments” of alcohol.

Oftentimes there are possible interactions between the covariates and the treatments. For example, if we were trying to determine if two different brands of fertilizer resulted in the same crop yield or not and we were measuring insect infestation as our covariate. We would expect that the greater the infestation, the less crop yield. Beyond this, it could be that one of the fertilizers contained something that insects enjoyed eating and thus, would be more susceptible more substantial infestations. It would then be difficult to tell if low crop yield for this fertilizer would be result of the fertilizer simply not working or due to the fact that the crop was eaten by the insects.

Thus, or analysis of covariance begins much the same way as it did for the blocking designs and the factorial data. We start by testing for interactions between the covariate and treatments. Using our example of two fertilizers and insect infestation, the model that leaves interaction as a possibility is,

$$y = \beta_1 \text{Fertilizer}_A + \beta_2 \text{Fertilizer}_B + \beta_3 \text{Infestation}_A + \beta_4 \text{Infestation}_B,$$

where Fertilizer A and B are dummy 0/1 indicator variables exactly like the ANOVA model and Infestation A and B are infestation values measured for the different treatments (infestation is scored from 0 to 10, with 0 being no infestation and 10 being the most infestation).

Given 8 measurements for each treatment, the re-



sponse vector and the design matrix are:<sup>26</sup>

$$\mathbf{Y} = \begin{bmatrix} 18 \\ 15 \\ 12 \\ 11 \\ 13 \\ 17 \\ 12 \\ 16 \\ 9 \\ 10 \\ 12 \\ 13 \\ 15 \\ 15 \\ 11 \\ 9 \end{bmatrix} \quad \mathbf{X} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 5 & 0 \\ 1 & 0 & 9 & 0 \\ 1 & 0 & 8 & 0 \\ 1 & 0 & 7 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 10 & 0 \\ 1 & 0 & 3 & 0 \\ 0 & 1 & 0 & 9 \\ 0 & 1 & 0 & 8 \\ 0 & 1 & 0 & 4 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 6 \\ 0 & 1 & 0 & 9 \end{bmatrix}$$

We can test for interaction by testing whether the Infestation slopes are the same for the different fertilizers. That is,

$$\mathbf{C} = [ 0 \quad 0 \quad 1 \quad -1 ],$$

and our general linear models program give us the results:

```
> general_linear (x, y, c);
beta =

    17.86079
    14.48320
    -0.67178
    -0.59096

f_test = 0.56192
p_value = 0.46793
t_test = 0.74961
```

Here the p-value is much larger than our 0.05 cut off, so that we will fail to reject our hypothesis that the slopes for the covariate given different treatments are different and thus, we can assume that there is no interaction between Infestation and Fertilizer.

Since we can ignore interaction, we can use a simpler model (one that leaves out the possibility for interaction) to test whether or not the fertilizers have the same effect on crop yield. This model simply

---

<sup>26</sup>The data given here were stolen from David Dickey and Jimmy Joi's web page: [http://www.stat.ncsu.edu/~st512/info/dickey/crsnotes/notes\\_5.htm](http://www.stat.ncsu.edu/~st512/info/dickey/crsnotes/notes_5.htm).

lumps the two Infestation variables into one. We can do this because we failed to detect any difference between the two variables in the more complex model. Thus, our new model is:

$$y = \beta_1 \text{Fertilizer}_A + \beta_2 \text{Fertilizer}_B + \beta_3 \text{Infestation},$$

and the design matrix becomes:

$$\mathbf{X} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 5 \\ 1 & 0 & 9 \\ 1 & 0 & 8 \\ 1 & 0 & 7 \\ 1 & 0 & 1 \\ 1 & 0 & 10 \\ 1 & 0 & 3 \\ 0 & 1 & 9 \\ 0 & 1 & 8 \\ 0 & 1 & 4 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 1 & 6 \\ 0 & 1 & 9 \end{bmatrix}.$$

The contrast matrix for determining if there is a difference between the treatments is

$$\mathbf{C} = [ 1 \quad -1 \quad 0 ],$$

and our program gives us the results:

```
> general_linear (x, y, c);
beta =

    14.48603
    11.51397
    -0.62940

f_test = 60.644
p_value = 2.9992e-06
t_test = 7.7874
```

Such a small p-value causes us to reject the hypothesis that both Fertilizer treatments are the same.

**NOTE:** This second test, where we have lumped the covariate into a single variable for all of the treatments in the model and are testing for differences between the treatments, is equivalent to conducting the test under the *adjusted treatment means* hypothesis that is sometimes mentioned in other statistics texts.

||

### Example 3.13.5.12 (Categorical Data)

Ever since 1969, linear models have been used to analyze categorical data [?]. ||

### 3.13.6 Linear Models with Multiple Dependent Variables

Suppose the observations, or dependent variables,  $y_{is}$ , are vectors with  $q$  correlated characteristics instead of single variables, as would be the case of multiple observations made on the same individual. A random sample of  $n$  of these vectors could be arranged in a rectangular array to form an  $n \times q$  matrix  $\mathbf{Y}$ , where the first row of  $\mathbf{Y}$  is the vector of characteristics observed on the first individual, the second row is the vector observed on the second individual and so on.

Assuming that the  $q$ -dimensional observation vector has a multivariate normal distribution, and that the  $n$  observations vectors are independent, we can extend the univariate model developed in Sections ?? through ?? to encompass the  $q$  correlated variables. The model now appears as:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad (3.13.24)$$

which looks exactly like the univariate general linear model in Equation ??, except in this case,  $\mathbf{Y}$  is an  $n \times q$  matrix and  $\boldsymbol{\beta}$  is an  $m \times q$  matrix. The matrix  $\mathbf{X}$ , the design matrix, is the same matrix of known constants that appeared in the univariate model. The hypothesis can be generalized to:

$$\mathbf{C}\boldsymbol{\beta}\mathbf{U} = \boldsymbol{\theta}, \quad (3.13.25)$$

where  $\mathbf{C}$  is an  $t \times p$  matrix  $\mathbf{U}$  is a  $q \times u$  matrix, and  $\boldsymbol{\theta}$  is an  $t \times u$  matrix and  $\mathbf{C}$  and  $\mathbf{U}$  are arbitrary matrices designed to yield the appropriate hypothesis.

#### Example 3.13.6.1 (Multiple Regression)

A series of animals were studied where cardiac output and mean blood pressure were measured while heart rate and respiration were varied. The data from this study can be found in Appendix ?. We will model with data with the formula:

$$(y_{i,1}, y_{i,2}) = \beta_0 + x_{i,1}\beta_1 + x_{i,2}\beta_2 \quad (3.13.26)$$

where

$y_{i,1}$  = mean blood pressure of the  $i$ -th animal,

$y_{i,2}$  = cardiac output of the  $i$ -th animal,

$x_{i,1}$  = respiration rate of the  $i$ -th animal,

$x_{i,2}$  = heart rate of the  $i$ -th animal.

Thus,

$$\mathbf{Y} = \begin{bmatrix} y_{1,1} & y_{1,2} \\ \vdots & \vdots \\ y_{n,1} & y_{n,2} \end{bmatrix}, \quad \mathbf{X} = \begin{bmatrix} 1 & x_{1,1} & x_{1,2} \\ \vdots & \vdots & \vdots \\ 1 & x_{n,1} & x_{n,2} \end{bmatrix},$$

and

$$\boldsymbol{\beta} = \begin{bmatrix} \beta_{0,1} & \beta_{0,2} \\ \beta_{1,1} & \beta_{1,2} \\ \beta_{n,1} & \beta_{n,2} \end{bmatrix}.$$

Some questions that we might ask about this data are

1. Does respiration rate affect cardiac output and mean blood pressure?
2. Does heart rate affect cardiac output and mean blood pressure?

To answer the first question, we test the hypothesis that the respiration rate regression coefficients are zero:

$$H_0 : \beta_{1,1} = \beta_{1,2} = 0.$$

We can convert this hypothesis into matrix form using Equation ?? by defining  $\mathbf{C}$  and  $\mathbf{U}$  such that

$$\mathbf{C} = [ 0 \quad 1 \quad 0 ] \quad \text{and} \quad \mathbf{U} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

which yields:

$$\begin{aligned} [ 0 \quad 1 \quad 0 ] \begin{bmatrix} \beta_{0,1} & \beta_{0,2} \\ \beta_{1,1} & \beta_{1,2} \\ \beta_{n,1} & \beta_{n,2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} &= [ 0 \quad 0 ] \\ [ \beta_{1,1} & \beta_{1,2} ] \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} &= [ 0 \quad 0 ] \\ [ \beta_{1,1} & \beta_{1,2} ] &= [ 0 \quad 0 ] \end{aligned}$$

To answer the second question, we test the hypothesis that the heart rate coefficients are zero:

$$H_0 : \beta_{2,1} = \beta_{2,2} = 0,$$

for which

$$\mathbf{C} = [ 0 \quad 0 \quad 1 ] \quad \text{and} \quad \mathbf{U} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Since, in general, these two tests will not be independent, we should make them simultaneously. To do this, let

$$\mathbf{C} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathbf{U} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Using Equation ??, this yields:

$$\mathbf{C}\beta\mathbf{U} = \begin{bmatrix} \beta_{1,1} & \beta_{1,2} \\ \beta_{n,1} & \beta_{n,2} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

We are now ready to use the multivariable version of the general linear models program found in Appendix ||

### Example 3.13.6.2 (Hotelling T<sup>2</sup>)

Hotelling T<sup>2</sup> tests allow us to run simultaneous paired *t*-tests on *q* pairs of characteristics. Consider, in this case, that we have measured mean blood flow, mean blood pressure, cerebro-vascular resistance in a series of experimental subjects both before and after the administration of epinephrine. Using a Hotelling T<sup>2</sup> we can answer the question: Did the drug change the blood flow, pressure and resistance significantly?

The model for this experiment is thus

$$\mathbf{Y} = \beta\mathbf{X}$$

where

$y_{i,1}$  = blood flow before

$y_{i,2}$  = blood pressure before

$y_{i,3}$  = resistance before

$y_{i,4}$  = blood flow after

$y_{i,5}$  = blood pressure after

$y_{i,6}$  = resistance after

$$\beta = [ \beta_{fb} \quad \beta_{pb} \quad \beta_{rb} \quad \beta_{fa} \quad \beta_{pa} \quad \beta_{ra} ] \text{ and } \begin{bmatrix} 1 \\ 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix},$$

where the subscript 'b' refers to a measurement taken before the treatment and the subscript 'a' refers to a measurement taken after.

To answer our question about whether the drug changed blood flow, pressure and resistance, we ask

if the parameters for before and after measurements have changed. Thus, we form the hypothesis:

$$\begin{aligned} H_0 : \beta_{fb} - \beta_{fa} &= 0 \\ \beta_{pb} - \beta_{pa} &= 0 \\ \beta_{rb} - \beta_{ra} &= 0 \end{aligned}$$

This leads us to define **C** and **U** as:

$$\mathbf{C} = [1] \text{ and } \mathbf{U} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix},$$

so that

$$\mathbf{C}\beta\mathbf{U} = [ 0 \quad 0 \quad 0 ],$$

or, in other words,

$$\begin{bmatrix} (\beta_{bf} - \beta_{ba}) & (\beta_{pf} - \beta_{pa}) & (\beta_{rf} - \beta_{ra}) \end{bmatrix} = [ 0 \quad 0 \quad 0 ].$$

The output from our program is... ||

### Example 3.13.6.3 (Multivariate ANOVA)

A series of twenty-four animals were studied by dividing them into six groups according to their diet and sex. The cardiac output, heart rate, and initial body weight of the animals were measured. Since body weight was thought to affect the level of response, it is considered a covariate. Cardiac output and heart rate are both dependent variables. Our model is thus,

$$\begin{aligned} (y_{i,1}, y_{i,2}) &= x_{i,1}\beta_1 + x_{i,2}\beta_2 + x_{i,3}\beta_3 \\ &\quad + x_{i,4}\beta_4 + x_{i,5}\beta_5 + x_{i,6}\beta_6 + x_{i,c}, \end{aligned}$$

where

$$x_{i,t} \begin{cases} 1 & \text{if animal received treatment } t, \\ 0 & \text{if animal did not receive treatment } t, \end{cases}$$

and

$x_{i,c}$  = initial body weight.

To test if cardiac output and heart rate vary with sex, we construct the contrast matrices:

$$\mathbf{C} = [ 1 \quad 1 \quad 1 \quad -1 \quad -1 \quad -1 \quad 0 ]$$

and

$$\mathbf{U} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Thus, our hypothesis is<sup>27</sup>:

$$\begin{bmatrix} (\beta_{1,1} + \beta_{2,1} + \beta_{3,1}) - (\beta_{4,1} + \beta_{5,1} + \beta_{6,1}) \\ (\beta_{1,2} + \beta_{2,2} + \beta_{3,2}) - (\beta_{4,2} + \beta_{5,2} + \beta_{6,2}) \end{bmatrix} = 0.$$

||

### 3.13.7 Conclusion

In conclusion, when viewing linear models from a matrix perspective, you build the model to mimic the physics or physiology that you believe to be operational - including terms that dominate the processes - and you're not stuck with a recipe book of 1-way ANOVA, ANCOVA etc. The matrix approach to linear modeling simply suppresses the mechanics of the computational process and lets you focus on the science - the model of what you think is going on (the process).

---

<sup>27</sup>See Appendix ?? for the derivation

# Chapter 4

## Basic Electronics

### 4.1 Symbols and Abbreviations

Quantity	Symbol	Unit	Abbreviation
Current	I	Ampere	A
Voltage	E or V	Volt	V
Resistance	R	Ohm	$\Omega$
Conductance	G	Siemens	S
Power	P	Watt	W
Capacitance	C	Farad	F
Inductance	L	Henry	H
Reactance	X	Ohm	$\Omega$
Impedance	Z	Ohm	$\Omega$
Gain	A		

$$\begin{aligned} \text{Angular Velocity} &= 2\pi\text{Hz} \\ &= \omega\text{Hz} \end{aligned}$$

where  $\omega$  (lowercase omega) is equal to  $2\pi$ .

**Anode:** Positive.

**Bel:**

$$B = \log_{10} \frac{P_{\text{output}}}{P_{\text{input}}}$$

**Capacitance:** The measure of a capacitor's ability to store energy for a given amount of voltage drop. Capacitors store energy in an electric field.

When a capacitor is faced with an increasing voltage, it acts as a load: drawing current as it absorbs energy (current going in the negative side and out the positive side, like a resistor). This energy is stored in the form of a field. As voltage increases, the field's force increases.

When a capacitor is faced with a decreasing voltage, it acts as a source: supplying current as it releases stored energy (current going out the negative side and in the positive side, like a battery).

$$C = \frac{i}{\frac{dv}{dt}}$$

or

$$i = C \frac{dv}{dt}$$

where  $i$  is the instantaneous current through the capacitor and  $v$  is the instantaneous voltage.

**Cathode:** Negative.

**Conductance:** The inverse of resistance.

$$\text{conductance} = \frac{1}{\text{resistance}}$$

**Current:** The actual movement of electrons between two points. In a closed circuit, the current is constant

### 4.2 Measurement Prefixes

Prefix	Abbreviation	Meaning
tera	T	$10^{12}$
giga	G	$10^9$
mega	M	$10^6$
kilo	K	$10^3$
milli	m	$10^{-3}$
micro	$\mu$	$10^{-6}$
nano	n	$10^{-9}$
pico	p	$10^{-12}$

### 4.3 Terminology and Definitions

1

**Angular Velocity:** The radians per second traveled by an AC current. This is primarily a means to express the frequency of an AC circuit in units of radians per second instead of cycles per second.

<sup>1</sup>Most of the definitions in this section were plagiarized from Tony R. Kuphaldt's *Lessons in Electric Circuits*

throughout and is not relative to two different points. Current is caused by voltage.

A useful analogy is to consider a several water reservoirs that are connected and are at different heights. The potential energy (voltage) measured between the highest reservoir and the lowest reservoir is going to be the greatest when compared to the potential energy (voltage) measured between the highest and second highest reservoirs or any other combination of reservoirs. However, the amount of water (current) flowing between each tank is going to be constant (this is assuming that there is a means to push water back to the highest reservoir). Even if the pipes between tanks are different sizes (and thus, potentially allowing more water to leave some tanks than others), an equilibrium will be reached before too long (assuming water is not being added from some external source). The amount of water flow (current) is limited by the smallest pipe because, in the long run, it determines how quickly the other tanks will be re-supplied with water.

Current is the rate at which electric charges move through a conductor.

**decibel:**

$$\text{dB} = 10 \log_{10} \frac{P_{\text{output}}}{P_{\text{input}}}$$

**Field Flux:** The total quantity, or effect, of a field through space. The flux of a field is roughly analogous to the current in a circuit.

**Field Force:** The amount of “push” that a field exerts over a certain distance, thus causing flux to form in space. The force of a field is roughly analogous to the voltage in a circuit.

**Gain:**

$$A = \frac{P_{\text{output}}}{P_{\text{input}}}$$

**Impedance:** The vector combination (summation) of resistance and reactance in a circuit.

**Inductor Reactance:** Opposition to changes in current that results in a phase shift and no dissipation in power. A term used with inductors. A resistor opposes changes in current, but does not induce a phase shift and does dissipate power (in the form of heat).

$$X = 2\pi fL$$

where  $f$  is the frequency of the alternating current.

**Joule’s Law:**  $P = I^2R$  (substituting Ohm’s law into the Power equation.)

**Kirchhoff’s Current Law:** The algebraic sum of all currents entering and exiting a node must equal zero.

**Kirchhoff’s Voltage Law:** The algebraic sum of all voltages in a loop must equal zero.

**Ohm’s Law (AC):**  $E = IZ$  (All quantities should be expressed in complex form).

**Ohm’s Law (DC):**  $E = IR$ .

**Polarity:** The +/- orientation of a voltage drop.

**Power:**  $P = IE$  (measured in Watts)

**Rectification:** Converting alternating current to direct current. Rectifier circuits are usually created with diodes.

**Resistance:** The measure of friction a component presents to the flow of electrons through it. In series circuits, the total resistance is the sum of all of the resistors. In parallel circuits, the total resistance is equal to

$$\Omega_T = \frac{1}{\frac{1}{\Omega_1} + \frac{1}{\Omega_2} \cdots + \frac{1}{\Omega_n}}$$

**Resistors:** These are rated both for the amount of resistance that they can insert into a circuit, but also how much heat energy that they can dissipate without sustaining damage. Resistance, by its nature, causes a dissipation of energy and this most often is in the form of heat. Resistance is measured in Ohms,  $\Omega$ , and the heat energy is measured as Watts,  $W$ . You can use the definition of power to derive the amount of energy dissipation that will be caused by inserting a resistor into a circuit and use that to determine what kind of resistor to use.

**Voltage:** The amount of potential energy available (work to be done) per unit charge to move electrons through a conductor. Because the concept of “potential energy” is relative to the source and the destination, voltage is always expressed as a quantity between two points. Often times the voltage measured between two points is called a “voltage drop”.

An example of voltage as potential energy is a battery that is not connected to anything. Even though there is no flow of electrons between the negative and the positive terminals of the battery (since neither terminal is connected to anything), there is still a voltage between the two poles.

Voltage is the specific work (or potential energy) per unit charge.

**Inductance:** The measure of an inductor’s ability to store energy for a given amount of current flow. Inductors store energy in a magnetic field.

When an inductor is faced with an increasing current, it acts as a load: dropping voltage as it absorbs energy (negative on the current entry side and positive on the current exit side, like a resistor).

When an inductor is faced with a decreasing current, it acts as a source: creating voltage as it releases stored energy (positive on the current entry side and negative on the current exit side, like a battery).

$$L = \frac{v}{\frac{di}{dt}}$$

or

$$v = L \frac{di}{dt}$$

where  $i$  is the instantaneous current through the capacitor and  $v$  is the instantaneous voltage.

Thus, the faster the current is changing, the larger amount of voltage will be dropped over the inductor.

# Chapter 5

## Basic Linear Systems




Figure 5.1.1: A simple circuit with a single resistor and a single inductor

### 5.1 Resistor and Inductor: First Order Differential Equations

#### 5.1.1 General Overview

An example of a simple circuit with a single resistor and a single inductor can be seen in Figure ???. The voltage drops of the various components and how to analyze them is listed here:

- Voltage Drop (Ohm's Law) for Resistors:  $E = IR$ , where  $E$  is the voltage, measured in volts (V),  $I$  is the current, measured in amperes (A), and  $R$  is the resistance, measured in ohms,  $\Omega$ .
- Voltage Drop for Inductors:  $E = L \frac{dI}{dt}$ , where  $L$  is the inductance, measured in henries (H).
- Kirchoff's Voltage Law (total voltage is equal to the sum of the voltage drops):  $E = IR + L \frac{dI}{dt}$ .

#### 5.1.2 DC

If  $R = 12\Omega$ ,  $L = 4H$  and  $E = 60V$ , then we can solve for  $I(t)$  using an integrating factor (as explained in Section ??),

$$\begin{aligned}
 12I + 4 \frac{dI}{dt} &= 60 \\
 3I + \frac{dI}{dt} &= 15 \\
 e^{3t} 3I + e^{3t} \frac{dI}{dt} &= 15e^{3t} \\
 \frac{d}{dt} [e^{3t} I] &= 15e^{3t} \\
 \int \frac{d}{dt} [e^{3t} I] &= \int 15e^{3t} \\
 e^{3t} I &= 5e^{3t} + C \\
 I(t) &= 5 + Ce^{-3t}.
 \end{aligned}$$

If  $I(0) = 0$ , then

$$\begin{aligned}
 0 &= 5 + C \\
 C &= -5
 \end{aligned}$$

$$I(t) = 5 - 5e^{-3t}$$

A graph of the current can be seen in Figure ??.






Figure 5.1.2: A graph of how the current increases to 5 amperes over time. Here  $E = 60$ .

### 5.1.3 AC

If  $E(t) = 60 \sin(30t)V$ , then we can solve for  $I(t)$  using an integrating factor:

$$\begin{aligned}
 12I + 4 \frac{dI}{dt} &= 60 \sin(30t) \\
 3I + \frac{dI}{dt} &= 15 \sin(30t) \\
 e^{3t} 3I + e^{3t} \frac{dI}{dt} &= 15e^{3t} \sin(30t) \\
 \frac{d}{dt} [e^{3t} I] &= 15e^{3t} \sin(30t) \\
 \int \frac{d}{dt} [e^{3t} I] &= \int 15e^{3t} \sin(30t) \\
 e^{3t} I &= 15 \left[ \frac{e^{3t}}{3^2 + 30^2} (3 \sin(30t) - 30 \cos(30t)) + C \right] \\
 I(t) &= \frac{5}{101} [\sin(30t) - 10 \cos(30t)] + C e^{-3t}.
 \end{aligned}$$

If  $I(0) = 0$ , then

$$\begin{aligned}
 0 &= -\frac{50}{101} + C \\
 C &= \frac{50}{101}
 \end{aligned}$$

and thus,

$$I(t) = \frac{5}{101} [\sin(30t) - 10 \cos(30t)] + \frac{50}{101} e^{-3t}$$

A graph of the current can be seen in Figure ??.




Figure 5.1.3: A graph of the current over time. Here  $E = 60 \sin(30t)$ . The green line shows a reference to input  $(1/2) \sin(30t)$ . Notice how the current is 90 degrees phase shifted from the input.

## 5.2 Resistor and Capacitor: First Order Differential Equations

### 5.2.1 General Overview

An example of a simple circuit with a single capacitor and a single resistor can be seen in Figure ??.

- Voltage Drop (Ohm's Law) for Resistors:  $E = IR$ , where  $E$  is the voltage, measured in volts (V),  $I$  is the current, measured in amperes (A), and  $R$  is the resistance, measured in ohms,  $\Omega$ .
- Voltage Drop for Capacitors:  $E = \frac{Q}{C}$ , where  $Q$  is the charge, measured in coulombs, and  $C$  is the capacitance, measured in farads (F).
- **NOTE:** Current is the change in charge over time, thus  $I = \frac{dQ}{dt}$ .
- Kirchoff's Voltage Law (total voltage is equal to the sum of the voltage drops):  $E = IR + \frac{Q}{C} = R \frac{dQ}{dt} + \frac{Q}{C}$




Figure 5.2.1: A simple circuit with a single capacitor and a single resistor

### 5.2.2 DC

If  $R = 5\Omega$ ,  $C = 0.05\text{F}$ , and  $E = 60\text{V}$ , then we can solve for  $Q(t)$  using an integrating factor,

$$\begin{aligned} 5 \frac{dQ}{dt} + \frac{1}{0.05}Q &= 60 \\ 5 \frac{dQ}{dt} + 20Q &= 60 \\ \frac{dQ}{dt} + 4Q &= 12 \\ e^{4t} \frac{dQ}{dt} + 4Qe^{4t} &= 12e^{4t} \\ \frac{d}{dt} [Qe^{4t}] &= 12e^{4t} \\ \int \frac{d}{dt} [Qe^{4t}] &= \int 12e^{4t} \\ Qe^{4t} &= 3e^{4t} + C \\ Q(t) &= 3 + Ce^{-4t}. \end{aligned}$$

If  $Q(0) = 0$ , then

$$\begin{aligned} 0 &= 3 + C \\ C &= -3 \end{aligned}$$

and thus,

$$Q(t) = 3 - 3e^{-4t} \quad (5.2.1)$$

To derive an equation for the current,  $I(t)$ , we simply take the derivative of Equation ??.

$$I(t) = \frac{dQ}{dt} = 12e^{-4t}.$$




Figure 5.2.2: A graph of how the current,  $I(t)$ , goes to zero as the capacitor builds its charge. Here  $E = 60$ .

A graph of the charge,  $Q(t)$ , building up and the current,  $I(t)$ , going to zero can be seen in Figure ??.

### 5.2.3 AC

If  $R = 2\Omega$ ,  $C = 0.01\text{F}$  and  $E(t) = 10 \sin(60t)$ , then we can solve for  $Q(t)$  using an integrating factor:

$$\begin{aligned} 2 \frac{dQ}{dt} + \frac{1}{0.01}Q &= 10 \sin(60t) \\ 2 \frac{dQ}{dt} + 100Q &= 10 \sin(60t) \\ \frac{dQ}{dt} + 50Q &= 5 \sin(60t) \\ e^{50t} \frac{dQ}{dt} + 50Qe^{50t} &= 5e^{50t} \sin(60t) \\ \frac{d}{dt} [Qe^{50t}] &= 5e^{50t} \sin(60t) \\ \int \frac{d}{dt} [Qe^{50t}] &= \int 5e^{50t} \sin(60t). \end{aligned}$$

Integrating both sides then gives us a general solution for  $Q(t)$ :

$$\begin{aligned} Qe^{50t} &= 5 \left[ \frac{e^{50t}}{50^2 + 60^2} (50 \sin(60t) - 60 \cos(60t)) + C \right] \\ Q(t) &= \frac{5}{2500 + 3600} [50 \sin(60t) - 60 \cos(60t)] \\ &\quad + Ce^{-50t} \\ Q(t) &= \frac{5}{122} \sin(60t) - \frac{6}{122} \cos(60t) + Ce^{-50t}. \end{aligned}$$




Figure 5.2.3: A graph of the charge,  $Q(t)$ , and current,  $I(t)$ . Here  $E = 10 \sin(60t)$




Figure 5.3.1: A simple circuit with a single resistor, a single inductor and a single capacitor.

If  $Q(0) = 0$ , then

$$0 = 0 - \frac{6}{122} + C$$

$$C = \frac{6}{122}$$

and thus,

$$Q(t) = \frac{5}{122} \sin(60t) - \frac{6}{122} \cos(60t) + \frac{6}{122} e^{-50t} \quad (5.2.2)$$

To derive an equation for the current,  $I(t)$ , we simply take the derivative of Equation ??.

$$I(t) = \frac{dQ}{dt} = \frac{300}{122} \cos(60t) + \frac{360}{122} \sin(60t) - \frac{300}{122} e^{-50t}$$

A graph of the charge and current can be seen in Figure ??.

## 5.3 Resistor, Inductor and Capacitor: Second Order Differential Equations

### 5.3.1 General Overview

An example of a simple circuit with a single resistor, a single inductor and a single capacitor can be seen in Figure ??.

- Voltage Drop (Ohm's Law) for Resistors:  $E = IR$ , where  $E$  is the voltage, measured in volts (V),  $I$  is the current, measured in amperes (A), and  $R$  is the resistance, measured in ohms,  $\Omega$ .

- Voltage Drop for Inductors:  $E = L \frac{dI}{dt}$ , where  $L$  is the inductance, measured in henries (H).
- Voltage Drop for Capacitors:  $E = \frac{Q}{C}$ , where  $Q$  is the charge, measured in coulombs, and  $C$  is the capacitance, measured in farads (F).
- **NOTE:** Current is the change in charge over time, thus  $I = \frac{dQ}{dt}$ .
- Kirchoff's Voltage Law (total voltage is equal to the sum of the voltage drops):  $E = IR + L \frac{dI}{dt} + \frac{Q}{C} = R \frac{dQ}{dt} + L \frac{d^2Q}{dt^2} + \frac{Q}{C}$

### 5.3.2 DC

If  $R = 20\omega$ ,  $L = 1\text{H}$ ,  $C = 0.002\text{F}$  and  $E = 12\text{V}$ , then

$$20 \frac{dQ}{dt} + \frac{d^2Q}{dt^2} + \frac{1}{0.002} Q = 12$$

$$\frac{d^2Q}{dt^2} + 20 \frac{dQ}{dt} + 500Q = 12.$$

We can use the method of undetermined coefficients (explained in Section ??) to solve for the charge,  $Q(t)$ , and by differentiating  $Q(t)$ , the current,  $I(t)$ .

First, we must solve for  $Q_c(t)$ . Here the complementary equation is:

$$\frac{d^2Q}{dt^2} + 20 \frac{dQ}{dt} + 500Q = 0,$$

and the characteristic equation is:

$$r^2 + 20r + 500 = 0.$$

The roots of the characteristic equation are  $r_1 = -10 + 20i$  and  $r_2 = -10 - 20i$  and thus,

$$Q_c(t) = e^{-10t} (c_1 \sin(20t) + c_2 \cos(20t)).$$

To solve for  $Q_p(t)$  we begin by setting it to an unknown constant since  $E(t) = 12$ . Thus,

$$\begin{aligned} Q_p(t) &= A \\ Q_p'(t) &= 0 \\ Q_p''(t) &= 0 \end{aligned}$$

and we substitute these into the original differential equation:

$$\begin{aligned} 500A &= 12 \\ A &= \frac{3}{125}, \end{aligned}$$

thus,

$$Q_p(t) = \frac{3}{125}.$$

Combining  $Q_c(t)$  and  $Q_p(t)$  gives us a general solution to the differential equation:

$$Q(t) = e^{-10t} (c_1 \sin(20t) + c_2 \cos(20t)) + \frac{3}{125}.$$

If the initial charge is zero, that is  $Q(0) = 0$ , then

$$\begin{aligned} c_2 + \frac{3}{125} &= 0 \\ c_2 &= -\frac{3}{125}. \end{aligned}$$

If the initial current is also zero, that is  $I(t) = 0$ , then we must first differentiate  $Q(t)$ , since  $I(t) = Q'(t)$ .

$$\begin{aligned} I(t) = \frac{dQ}{dt} &= -10e^{-10t} (c_1 \sin(20t) + c_2 \cos(20t)) \\ &\quad + e^{-10t} (20c_1 \cos(20t) - 20c_2 \sin(20t)), \end{aligned}$$

and then

$$\begin{aligned} I(0) = 0 &= -10c_2 + 20c_1 \\ &= \frac{30}{125} + 20c_1 \\ c_1 &= -\frac{30}{2500} = -\frac{3}{250} \end{aligned}$$

Thus, with  $c_1 = -\frac{3}{250}$  and  $c_2 = -\frac{3}{125}$ , the formula for the charge is:

$$Q(t) = -e^{-10t} \left( \frac{3}{250} \sin(20t) + \frac{3}{125} \cos(20t) \right) + \frac{3}{125},$$




Figure 5.3.2: A graph of the charge,  $Q(t)$ , and current,  $I(t)$ . Here  $E = 12V$

and the formula for the current is:

$$\begin{aligned} I(t) &= 10e^{-10t} \left[ \frac{3}{250} \sin(20t) + \frac{3}{125} \cos(20t) \right] \\ &\quad - e^{-10t} \left[ \frac{60}{250} \cos(20t) - \frac{60}{125} \sin(20t) \right] \\ &= e^{-10t} \left[ \frac{15}{125} \sin(20t) + \frac{60}{250} \cos(20t) \right] \\ &\quad - e^{-10t} \left[ \frac{60}{250} \cos(20t) - \frac{60}{125} \sin(20t) \right] \\ &= e^{-10t} \frac{15}{125} \sin(20t) + e^{10t} \frac{60}{125} \sin(20t) \\ &= \frac{3}{5} e^{-10t} \sin(20t). \end{aligned}$$

A graph of the charge and the current can be seen in Figure ??

### 5.3.3 AC

If  $E(t) = 12 \sin(10t)$ , then

$$\frac{d^2Q}{dt^2} + 20\frac{dQ}{dt} + 500Q = 12 \sin(10t).$$

We can use the method of undetermined coefficients to solve for  $Q(t)$  and  $I(t)$ .

In Section ?? we found the solution for the complementary equation,  $Q_c(t)$  to be:

$$Q_c(t) = e^{-10t} (c_1 \cos(20t) + c_2 \sin(20t)).$$

To solve for  $Q_p(t)$  we begin by setting it to an equa-

tion similar to the function for  $E(t)$ . Thus,

$$\begin{aligned} Q_p(t) &= A \cos(10t) + B \sin(10t) \\ Q'_p(t) &= -10A \sin(10t) + 10B \cos(10t) \\ Q''_p(t) &= -100A \cos(10t) - 100B \sin(10t) \end{aligned}$$

and we substitute these into the original differential equation:

$$\begin{aligned} &(-100A \cos(10t) - 100B \sin(10t)) \\ &\quad + 20(-10A \sin(10t) + 10B \cos(10t)) \\ &\quad + 500(A \cos(10t) + B \sin(10t)) = 12 \sin(10t) \end{aligned}$$

which reduces to

$$\begin{aligned} (2A + B) \cos(10t) + (-A + 2B) \sin(10t) \\ = \frac{12}{200} \sin(10t). \end{aligned}$$

This gives us the system of equations:

$$\begin{aligned} 2A + B &= 0 \\ -A + 2B &= \frac{12}{200}, \end{aligned}$$

and from these we determine that  $A = \frac{-3}{250}$  and  $B = \frac{3}{125}$ .

Combining  $Q_c(t)$  and  $Q_p(t)$  gives us a general solution to the differential equation:

$$\begin{aligned} Q(t) &= e^{-10t} (c_1 \cos(20t) + c_2 \sin(20t)) \\ &\quad - \frac{3}{250} \cos(10t) + \frac{3}{125} \sin(10t). \end{aligned}$$

If the initial charge is zero, that is  $Q(0) = 0$ , then

$$\begin{aligned} c_1 - \frac{3}{250} &= 0 \\ c_1 &= \frac{3}{250} \end{aligned}$$

If the initial current is also zero, that is  $I(t) = 0$ , then we must first differentiate  $Q(t)$ , since  $I(t) = Q'(t)$ .

$$\begin{aligned} I(t) &= \frac{dQ}{dt} = -10e^{-10t} (c_1 \cos(20t) + c_2 \sin(20t)) \\ &\quad + e^{-10t} (20c_1 \sin(20t) - 20c_2 \cos(20t)) \\ &\quad + \frac{30}{250} \sin(10t) + \frac{30}{125} \cos(10t), \end{aligned}$$




Figure 5.3.3: A graph of the charge,  $Q(t)$ , and current,  $I(t)$ . Here  $E = 12 \sin(10t)$

and then

$$\begin{aligned} I(0) &= \frac{-30}{250} + 20c_2 + \frac{30}{125} = 0 \\ c_2 &= \frac{-3}{500}. \end{aligned}$$

Thus, with  $c_1 = \frac{-3}{250}$  and  $c_2 = \frac{3}{500}$ , the formula for the charge is:

$$\begin{aligned} Q(t) &= e^{-10t} \left[ \frac{3}{250} \cos(20t) + \frac{-3}{500} \sin(20t) \right] \\ &\quad - \frac{3}{250} \cos(10t) + \frac{3}{125} \sin(10t), \end{aligned}$$

and the equation for the current is:

$$\begin{aligned} I(t) &= -10e^{-10t} \left[ \frac{3}{250} \cos(20t) + \frac{-3}{500} \sin(20t) \right] \\ &\quad + e^{-10t} \left[ \frac{60}{250} \sin(20t) - \frac{-60}{500} \cos(20t) \right] \\ &\quad + \frac{30}{250} \sin(10t) + \frac{30}{125} \cos(10t), \end{aligned}$$

which reduces to:

$$\begin{aligned} I(t) &= e^{-10t} \left[ \frac{-120}{500} \cos(20t) - \frac{45}{250} \sin(20t) \right] \\ &\quad + \frac{30}{250} \sin(10t) + \frac{30}{125} \cos(10t). \end{aligned}$$

A graph of the charge and the current can be seen in Figure ??

## 5.4 Operational Amplifiers (opamps)

### 5.4.1 General Overview

# Chapter 6

## Making the “Data Modeling: HOWTO”: HOWTO

### 6.1 Introduction

Here we are attempting to document how this HOWTO was created. This is useful primarily for us, the authors, as we have found it is hard to remember all of the various details of different graphic programs or even how to manipulate the CVS directory.

### 6.2 Drawing Graphs

#### 6.2.1 GNUPlot

Here’s how the 3-D graph used to illustrate Lagrange multipliers (see Figure ?? on page ??) was created:

```
gnuplot> set term postscript color
gnuplot> set output "graph3d.eps"
gnuplot> set hidden3d
gnuplot> set contour
gnuplot> splot [-10:10][-10:10][-200:200]\
> -(x**2 + y**2)-10, x**2 + y**3
```

The first command, `set term`, sets the output to be a postscript file (with color). The second command sets the name of the file. `set hidden3d` tells gnuplot to let surfaces in the foreground block the visibility of those that are behind them. To turn this option off, and return to a simply wire frame rendering of your graphs, use the command `set nohidden3d`. `set contour` tells gnuplot to draw the contour lines on the plane beneath the graphs. The final command creates the graph. If you leave off the first two commands, gnuplot will draw the graph to a new window instead of writing it to a file.

#### 6.2.2 Octave

One thing to note about using Octave, is that to generate graphs, it uses GNUPlot, so a lot of the commands are the same.

Here is how to create one of the normal curve graphs that are so common in the paper:

```
> x = linspace(-8, 8, 100);
> y = normal_pdf(x);
> plot(x, y)
```

If we wanted to replace the “line 1” label that is the default, we could change our plot command to add our own version:

```
> plot(x, y, ‘;mu = 0, sigma = 1;’)
```

To export the graph as a color EPS file, all you need to do is:

```
> gset term postscript eps color;
> gset output "normal_curve.eps"
> replot
```

Sometimes when you are using octave, you may want to export the data that you are working with so you can import it into another graphing program other than gnuplot (for example, you may want to import the data into xmgrace). To do this requires a little bit of a trick since you’ll notice that octave exports vectors as a single line with a whitespace between the elements. While this makes perfect sense to do, it is often the case that you have a vector representing the values for the x-axis and a vector (or vectors) representing values on the y-axis (as we do in the example we just gave). If we were to export these vectors directly, we’d get the values for the x

and the y axis on two different lines and this is not what we want. Instead, we'd like to have each line in our output file correspond to a value from the x-axis vector and a value from the y-axis vector (or values from y-axis vectors). To do this, you simply create a new matrix that is made up of the transposes of these vectors. To continue with the previous example, if we wanted to export our data in a format that could easily be imported into xmgrace we would do this:

```
> x = linspace(-8, 8, 100);
> y = normal_pdf (x);
> z = [x', y'];
> save -ascii 'output_file.txt' z
```

### 6.2.3 Octave With Grace

This is a step by step guide to how Figure ?? was created.

Get data defining the  $\chi_1^2$  curve from octave:

```
octave> x = linspace(0, 8, 100);
octave> y = chisquare_pdf (x, 1);
octave> plot(x, y) # verify the graph
octave> z = [x', y'];
octave> save -ascii "chi_square1.dat" z
octave> exit
```

Set the GRACE\_EDITOR environment variable: The default editor for Grace data files is Vi. If this is fine with you, you can skip this step. Otherwise, you can either set the variable on the command line:

```
shell> export GRACE_EDITOR emacs
```

or edit your shell configuration file. Editing your configuration file has the advantage of setting things up for your future sessions with grace as well.

Now import the data into Grace:

```
shell> xmgrace chi_square1.dat &
```

You can modify the X-axis and the Y-axis to display things the way you want them to look by choosing **Axis properties** from the **Plot** menu.

Now we want to highlight the area under the curve specified by the chi-square statistic. To do this, we first duplicate the data set. You can do this in a number of ways. Here we will choose **Data sets...** from the **Edit** menu. This will bring up a window listing the existing data sets. Click on the dataset to select it. Now right click on the dataset and choose **Duplicate** from the menu.

We will now edit the duplicate data set to remove all of the points that occur before our chi-square value. We do this by right clicking on the new data set and selecting **Edit→In text editor** from the menu. Now, delete all of the data rows that come before  $X = 5.76$ . After you have done this, save the changes to the data, close the editor and click on the **Accept** button at the bottom of the window.

To color the area the under second, smaller dataset we now select **Set appearance** from the **Plot** menu and

- select the second set in the box at the top of the window
- click on the **Line** tab
- select **To baseline** from **Type** menu
- select a nice color and pattern.
- click on the **Accept** button at the bottom of the window.

We can also add text and arrows to graph with drawing tools by selecting **Drawing objects** from the **Window** menu.

To create EPS and PDF files (EPS files are needed to create DVI output and PDF files are needed to create PDF output) select **Print setup** from the **File** menu. Now select **eps** or **pdf**<sup>1</sup> from the **Device** menu. Now click **Accept** at the bottom of the window. Now select **Print** from the **File** menu.

### 6.2.4 Phase Plots With Maple

The phase plot illustrations in the section on Differential Equations, Section ??, were drawn using maple. The first thing you need to do is declare that you want to use the DEtools. This is done with the command:

```
> with(DEtools):
```

To draw the plot in Figure ??, the following commands were then used:

```
> deq1 := D(y)(t) = (-2 * (y(t) - 4));
> plot1 := DEplot(deq1, y(t), \
t=-3..3, y=-3..3, arrows=LINE):
> plots[display]({plot1});
```

<sup>1</sup>xmgrace doesn't do the best job exporting PDF images so it is sometimes better to export an EPS image and use **epstopdf** to create the PDF version. You will just have to do it both ways and decide which looks better.



To draw the plot in Figure ??, the following commands were used:

```
> deq2 := D(y)(t) =\
(y(t) - 1)*(y(t) - 2)*(y(t) - 3);
> plot2 := DEplot(deq2, y(t),\
t=-2..2, y=0..4, arrows=LINE):
> plots[display]({plot2});
```

## 6.3 Creating PDF Files From EPS Files

In order for both the command line `latex` and `pdflatex` programs to work correctly with a single file, `datamodel.tex` all illustrations be in both PDF and EPS formats. This is because `xdvi` will only display EPS images and PDF viewers will only display images in the PDF format. If the program you are working with will not export both types of images, EPS images can be converted using `epstopdf`.

## 6.4 CVS

All of the source LaTeX files as well as illustrations are stored in a CVS repository: see

<http://sourceforge.net>

or specifically, follow the link to the `datamodel` website,

<http://sourceforge.net/projects/datamodel/>

and follow the link to CVS.

### 6.4.1 Obtaining the Source

As a viewer of the project, you can either download the entire project or view the project files and selectively download what you are interested in. The page:

[http://sourceforge.net/cvs/?group\\_id=44909](http://sourceforge.net/cvs/?group_id=44909)

gives detailed instructions for anonymous ftp of the project files.

As a developer, to access them you must first make sure that the environment variable `CVS_RSH` is set to `ssh`. You can do this with the command:

```
shell> export CVS_RSH=ssh
```

or by adding `export CVS_RSH=ssh` to your shell configuration file and reloading the shell. If you use the latter method you will, of course, only have to modify the file one time and never have to worry about this step again.

With the `CVS_RSH` environment variable set, the `datamodel` project can then be downloaded using the command:

```
shell> cvs -z3\
> -d:ext:username@cvs.sourceforge.net:\
> /cvsroot/datamodel co datamodel
```

where `username` is your `sourceforge.net` username.

### 6.4.2 Updating The CVS Repository

After you have modified the files, you can update the CVS repository using the commands:

```
shell> cvs update
shell> cvs commit
```

When you run the `cvs commit` command, you will be asked to write a comment about the changes you made in the editor of your choice (defined by the `EDITOR` environment variable, so if you don't want to use `Vi` for this, you had better set this).

### 6.4.3 Adding New Files

You can add new files to the CVS repository with the commands:

```
shell> cvs add new_file_name
shell> cvs commit
```

**NOTE:** Once you add a new file, you can not remove it from the CVS repository.

## 6.5 Setting up a CVS Repository From Scratch

While this was not necessary for this current project, it is such a useful thing to do that it deserves mention here.

The very first thing you need to do is make sure you have the necessary software installed on your computer. That is, make sure you have `cvs`, `rcs` and `ssh` installed. Typical RedHat installs take care of this for you.

## 6.5.1 Creating and Initializing the Repository

This section describes a creation and initialization stage that is only required the very first time you try to get CVS working on your computer.<sup>2</sup> Once this is done, you can add new projects or directories to the repository without having to go through this process.

You first need to create a directory on your computer that can be used as the repository. Usually this is set to `/home/cvsroot/` but you can put it anywhere you want. You will also need to give read and write access to any users or groups that will be using the repository. For example, you may wish to use the following commands:

```
shell> su
shell# export CVSROOT=/home/cvsroot
shell# groupadd cvs
shell# useradd -g cvs -d $CVSROOT cvs
shell# mkdir $CVSROOT
shell# chgrp -R cvs $CVSROOT
shell# chmod o-rwx $CVSROOT
shell# chmod u+rwx $CVSROOT
shell# chmod g+rwx $CVSROOT
shell# cvs init
```

Now add add users to the cvs group.

```
bash# usermod -G cvs some_username
```

where `some_username` is a user you want to have read and write access to the repository.

## 6.5.2 User Environment Variables

There are three main environment variables you will need to have set, `EDITOR`, `CVS_RSH` and `CVSROOT`. Chances are, `EDITOR` is already set the way you like it. `CVS_RSH` must be set to `ssh`. `CVSROOT` is the location of the directory that will contain the repository. If it is not, now is a good chance to take care of it. It is best to set these variables in your shell configuration file. For example, adding the following lines:

```
export CVS_RSH=ssh
export EDITOR=emacs
export CVSROOT=/home/cvsroot
```

---

<sup>2</sup>A lot of the material in this section was plagiarized from the web page: `CVS-RCS-HOWTO.html`, written by Alavoor Vasudevan.

## 6.5.3 Adding a Project to the Repository

Now that you have the CVS repository initialized and you have your own environment variables set, you can add a project to it. To do this, first change to the project directory:

```
shell> cd $HOME/my_project_dir/
```

and import the directory:

```
cvs import my_project_dir vendor_1_0 ref_1_0
```

Here `vendor_1_0` and `ref_1_0` are just vendor and revision tags that are required for the initial import. I just leave them as is and things work just fine.

## 6.5.4 Setting up a Client Computer for CVS Access

Once you have your CVS repository running on your host computer, you may want to access it from other computers (clients). To do this, simply set the `EDITOR`, `CVS_RSH` and `CVSROOT` environment variables. In this case, however, you may only wish to set the `CVSROOT` environment variable on the command line as you may wish to use the computer to access CVS repositories on a variety of host machines. Thus, you may wish to add the following lines to your shell configuration file:

```
export CVS_RSH=ssh
export EDITOR=emacs
```

and, on the command line:

```
shell> export CVSROOT='':ext:\
> username@cvs_server_box.domain.edu:\
> /home/cvsroot''
```

Now you should be able to obtain the CVS projects stored on the host computer.

## 6.5.5 Building the Data Modeling Document

Due to the presence of both embedded references and an index, once you obtain the `datamodel` files, you must follow these steps to build the document:

```
shell> latex datamodel.tex
shell> makeindex datamodel.idx
shell> latex datamodle.tex
```

# Appendix A

## Octave Programs

### A.1 Univariate General Linear Models

```
function beta=general_linear(x,y,c)

# y is a column (nx1) vector of observations
# x is the design matrix (nxm)
# c is the contrast matrix
# NOTE: it is assumed that c is set up so that theta = 0

n = max(size(x)); # number of observations
p = min(size(x)); # number of parameters
num_tests = min(size(c)); # number of tests

beta = inv(x'*x)*x'*y # estimate the parameters

# test model
numerator = beta'*c'*inv((c*inv(x'*x)*c'))*c*beta / num_tests;
denominator = (y'*y - y'*x*beta)/(n-p); # this is also known as MSE
f_test = numerator / denominator
p_value = 1 - f_cdf(f_test, num_tests, n-p) # calculate p-value

if (num_tests == 1)
    t_test = sqrt(f_test)
end

# generate graphical output
MIN_F_DIST = 0;
MAX_F_DIST = 20;
NUM_STEPS = 100;

graph_x = linspace(MIN_F_DIST, MAX_F_DIST, NUM_STEPS);
graph_y = f_pdf(graph_x, num_tests, n-p);

description = sprintf("g;F(%d, %d)-Dist;", num_tests, n-p);
plot(graph_x, graph_y, description, f_test, 0, "r*;Your Data;")
```

```
endfunction
```

## A.2 Generating Normal(0,1) Random Variables

```
function normals=generate_normals(n)
  ## NOTE: if 'n' is odd, then you'll get floor(n) random variables
  for i = 1:(n/2)
    u1 = rand();
    u2 = rand();

    r = sqrt(-2*log(u1));
    theta = 2 * pi * u2;

    normals((i*2)-1) = r * cos(theta);
    normals(i*2) = r * sin(theta);
  endfor
endfunction
```

## A.3 Generating a Random Sample from a Discrete Distribution

```
function discretess=generate_discretess(n)

  ## First map the distribution to the interval [0,1].
  dist(1) = 0.25;
  dist(2) = 0.5;
  dist(3) = 0.25;

  for i = 1:(n)
    u1 = rand();

    total = 0;
    index = 1;
    for j = 1:length(dist)
      total = total + dist(j);
      if (u1 < total)
        index = j;
        break;
      endif
    endfor

    u2 = rand();

    discretess(i) = u2 + (index - 1);
  endfor
endfunction
```



# Appendix B

## Rather Large Matrices

### B.1

The design matrix used in Example ???. Only the first 31 of 60 rows are shown.

$$\mathbf{X} = \begin{bmatrix}
 \underbrace{1}_{\mu} & \underbrace{1 \ 0 \ 0 \ 0 \ 0}_{\text{Block}} & \underbrace{1}_{\text{Variation}} & \underbrace{1 \ 0 \ 0 \ 0 \ 0}_{\text{Block} \times \text{Variation}} & \underbrace{1 \ 0 \ 0 \ 0}_{\text{Spacing}} & \underbrace{1 \ 0 \ 0 \ 0}_{\text{Variation} \times \text{Spacing}} \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & 1 & 1 \ 0 \ 0 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & 1 & 1 \ 0 \ 0 \ 0 \ 0 & 0 \ 0 \ 1 \ 0 & 0 \ 0 \ 1 \ 0 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & 1 & 1 \ 0 \ 0 \ 0 \ 0 & 0 \ 0 \ 0 \ 1 & 0 \ 0 \ 0 \ 1 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & 1 & 1 \ 0 \ 0 \ 0 \ 0 & -1 \ -1 \ -1 \ -1 & -1 \ -1 \ -1 \ -1 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & -1 & -1 \ 0 \ 0 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 & -1 \ 0 \ 0 \ 0 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & -1 & -1 \ 0 \ 0 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 & 0 \ -1 \ 0 \ 0 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & -1 & -1 \ 0 \ 0 \ 0 \ 0 & 0 \ 0 \ 1 \ 0 & 0 \ 0 \ -1 \ 0 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & -1 & -1 \ 0 \ 0 \ 0 \ 0 & 0 \ 0 \ 0 \ 1 & 0 \ 0 \ 0 \ -1 \\
 1 & 1 \ 0 \ 0 \ 0 \ 0 & -1 & -1 \ 0 \ 0 \ 0 \ 0 & -1 \ -1 \ -1 \ -1 & 1 \ 1 \ 1 \ 1 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & 1 & 0 \ 1 \ 0 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & 1 & 0 \ 1 \ 0 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & 1 & 0 \ 1 \ 0 \ 0 \ 0 & 0 \ 0 \ 1 \ 0 & 0 \ 0 \ 1 \ 0 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & 1 & 0 \ 1 \ 0 \ 0 \ 0 & 0 \ 0 \ 0 \ 1 & 0 \ 0 \ 0 \ 1 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & -1 & 0 \ -1 \ 0 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 & -1 \ 0 \ 0 \ 0 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & -1 & 0 \ -1 \ 0 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 & 0 \ -1 \ 0 \ 0 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & -1 & 0 \ -1 \ 0 \ 0 \ 0 & 0 \ 0 \ 1 \ 0 & 0 \ 0 \ -1 \ 0 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & -1 & 0 \ -1 \ 0 \ 0 \ 0 & 0 \ 0 \ 0 \ 1 & 0 \ 0 \ 0 \ -1 \\
 1 & 0 \ 1 \ 0 \ 0 \ 0 & -1 & 0 \ -1 \ 0 \ 0 \ 0 & -1 \ -1 \ -1 \ -1 & 1 \ 1 \ 1 \ 1 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & 1 & 0 \ 0 \ 1 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & 1 & 0 \ 0 \ 1 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & 1 & 0 \ 0 \ 1 \ 0 \ 0 & 0 \ 0 \ 1 \ 0 & 0 \ 0 \ 1 \ 0 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & 1 & 0 \ 0 \ 1 \ 0 \ 0 & 0 \ 0 \ 0 \ 1 & 0 \ 0 \ 0 \ 1 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & 1 & 0 \ 0 \ 1 \ 0 \ 0 & -1 \ -1 \ -1 \ -1 & -1 \ -1 \ -1 \ -1 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & -1 & 0 \ 0 \ -1 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 & -1 \ 0 \ 0 \ 0 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & -1 & 0 \ 0 \ -1 \ 0 \ 0 & 0 \ 1 \ 0 \ 0 & 0 \ -1 \ 0 \ 0 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & -1 & 0 \ 0 \ -1 \ 0 \ 0 & 0 \ 0 \ 1 \ 0 & 0 \ 0 \ -1 \ 0 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & -1 & 0 \ 0 \ -1 \ 0 \ 0 & 0 \ 0 \ 0 \ 1 & 0 \ 0 \ 0 \ -1 \\
 1 & 0 \ 0 \ 1 \ 0 \ 0 & -1 & 0 \ 0 \ -1 \ 0 \ 0 & -1 \ -1 \ -1 \ -1 & 1 \ 1 \ 1 \ 1 \\
 1 & 0 \ 0 \ 0 \ 1 \ 0 & 1 & 0 \ 0 \ 0 \ 1 \ 0 & 1 \ 0 \ 0 \ 0 & 1 \ 0 \ 0 \ 0 \\
 \vdots & \vdots \ \vdots \ \vdots \ \vdots \ \vdots & \vdots & \vdots \ \vdots \ \vdots \ \vdots \ \vdots & \vdots \ \vdots \ \vdots \ \vdots & \vdots \ \vdots \ \vdots \ \vdots
 \end{bmatrix}$$

## B.2

To test for interaction between Variety and Spacing, the contrast matrix is:

$$\mathbf{C} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

## B.3

To test for a main effect for Spacing, the contrast matrix is:

$$\mathbf{C} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

## B.4

To test to see if the interaction term (Field×Variety) is significant, the contrast matrix is:

$$\mathbf{C} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

# Appendix C

## Data Used for Examples

### C.1

$y_1$ , Mean Blood Pressure	$y_2$ , Cardiac Output	$x_1$ , Respiration Rate	$x_2$ , Heart Rate
110.4	1.76	0.07	7.8
102.8	1.55	0.07	8.9
101.0	2.73	0.07	8.9
108.4	2.73	0.07	7.2
100.7	2.56	0.07	8.4
100.3	2.8	0.07	8.7
102.0	2.8	0.07	7.4
93.7	1.84	0.07	8.7
98.9	2.16	0.07	8.8
96.6	1.98	0.02	7.6
99.4	0.59	0.02	6.5
96.2	0.80	0.02	6.7
99.0	0.80	0.02	6.2
88.4	1.05	0.02	7.0
75.3	1.80	0.02	7.3
92.0	1.80	0.02	6.5
82.4	1.77	0.02	7.6
77.1	2.30	0.02	8.2
74.0	2.03	0.474	7.6
65.7	1.91	0.474	8.3
56.8	1.91	0.474	8.2
62.1	1.91	0.474	6.9
61.0	0.76	0.474	7.4
53.2	2.13	0.474	0.76
59.4	2.13	0.474	6.9
58.7	1.51	0.474	7.5
58.0	2.05	0.474	7.6



# Appendix D

## Derivations For the Curious

### D.1 Nernst-Planck Equation

When ions are in solution, there are three mechanisms for movement: brownian motion (thermal), ordered drift due to a potential (voltage) field, and diffusion, ordered drift down a concentration gradient.

Consider first, drift down a concentration gradient. In solution, each molecule is not stationary but is moving and the motion results in collisions with neighbors. In a region of high concentration, collisions are more likely than in regions of low concentration. Thus, at the interface between a high concentration and a low concentration, there will be a collision gradient, more collisions on the high concentration side than on the low concentration side (see figure ??). This gradient results in a drift of carriers into the low concentration region, increasing its concentration. An equilibrium is reached when the concentration equilibrates and the frequency of collision is spatially uniform. The flux associated with the drift is

$$j = -|Z|D \frac{d[C]}{dt} \tag{D.1.1}$$

where  $Z$  is the valence of the charge carrier and  $D$  is the diffusion constant.

Charge carriers are accelerated by the electrical attraction of the carrier within the electric field. As the charge is attracted, things get in the way that result in collisions. After each collision, velocity is lost resulting and is slowly recovered due to the acceleration caused by the attraction of the charge carrier and the potential field. To describe this, we start with the force that an unit charge feels within an electric field:

$$F = -qE = \frac{d(mv)}{dt} = \frac{mv_d}{\tau}$$

where  $F$  is the force,  $q$  is the unit charge and  $E$  is the electric field. Remember that the electric field,  $E = dV/dx$ , is simply the change in potential at a point. Now the attractive force will change the momentum of a charge carrier either positively (acceleration) or negatively (deceleration). We assume that the drift velocity is  $v_d$  and  $\tau$  is the time between the collisions of the charge carrier and something. From this, we can write the drift velocity between collisions due to the field (we ignore the collision events - acceleration and deceleration) as

$$v_d = -\frac{qE\tau}{m}$$

Now define the mobility of the charge,  $\mu$  as

$$\mu = \frac{q\tau}{m}$$

so that the drift velocity is

$$v_d = -\mu E$$

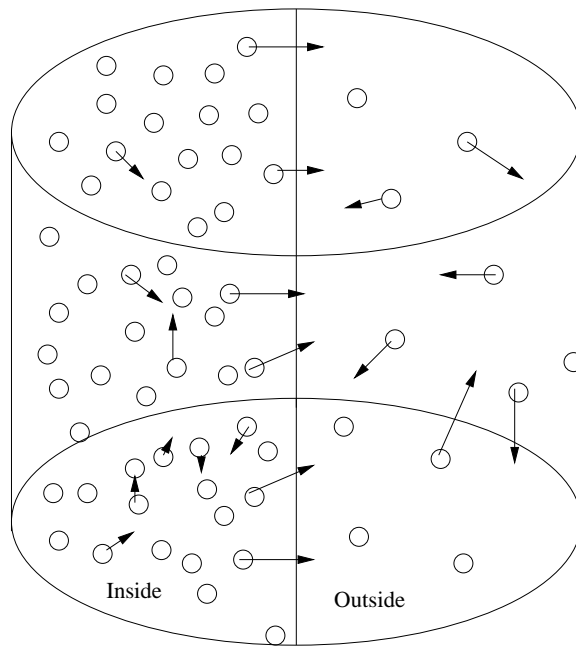


Figure D.1.1: A glass filled with high concentration Na on the left and low concentration Na on the right




Figure D.1.2: A + charge in solution attracted to the negative plate.

which simply states that in the presence of a spatially uniform electric field, the charge will move with a fixed velocity known as the drift velocity, that is proportional to the charge and inversely proportional to the mass of the charge.

Now the current density associated with the flow of charge within an electric field in a solution is:

$$j = v_d |Z| F [C]$$

where  $Z$  is the valence of the charge carrier,  $F$  is the number of coulombs of charge per mole of ion, and  $[C]$  is the concentration of charge carriers. But the drift velocity is  $-\mu E$  so that the current density can be written as

$$j = |Z| F [C] \mu E = \sigma E \quad \text{Ohm's law}$$

where  $\sigma = |Z| F \mu [C]$  is the conductivity. Note that the conductivity is proportional to the concentration of the ion carrier and its charge. Now, we combine the diffusive and electrical components of the flux under equilibrium conditions and have

$$-|Z| D \frac{d[C]}{dx} + |Z| F [C] \mu E = -D \frac{d[C]}{dx} + F [C] \mu \frac{dV}{dx} = 0$$

Now we integrate this across the interface (cell membrane) that separates the extracellular fluid from the cytoplasm and have

$$\begin{aligned} \int D \frac{d[C]}{dx} &= \int F [C] \mu \frac{dV}{dx} \\ \int D \frac{d[C]}{[C]} &= \int F \mu dV \\ D [\ln[C_{\text{out}}] - \ln[C_{\text{in}}]] &= F \mu (V_{\text{out}} - V_{\text{in}}) \\ V_{\text{out}} - V_{\text{in}} &= \frac{D}{F \mu} \ln \frac{[C_{\text{out}}]}{[C_{\text{in}}]} \end{aligned}$$

Einstein showed, in a cute little derivation that

$$D = \mu RT$$

where  $R$  is the gas constant (8.314 J/K mole at 27 C), and  $T$  is the absolute temperature. At 27 C,  $RT/F = 8.314 * 300 / 96487 = 25.8$  mV at 27 C. so for a monovalent cation or anion, the transmembrane potential due to a single charge carrier is:

$$V_{\text{membrane}}(\text{mV}) = \frac{RT}{F} \ln \frac{[C_{\text{out}}]}{[C_{\text{in}}]} = 25.8 \ln \frac{[C_{\text{out}}]}{[C_{\text{in}}]}$$

## D.2 Cubic Nonlinear ODE

$$\frac{du}{dt} = u(k - u^2) \quad (\text{D.2.1})$$

Solved using the method of separation of variables and some fancy algebra...

$$\begin{aligned} \frac{du}{dt} &= u(k - u^2) \\ \frac{du}{u(k - u^2)} &= dt \\ \int \frac{du}{u(k - u^2)} &= \int dt \end{aligned}$$

Now, we will break down  $\frac{1}{u(k-u^2)}$

First note:

$$(k - u^2) = (\sqrt{k} - u)(\sqrt{k} + u)$$

and thus,

$$\frac{1}{k - u^2} = \frac{1}{2\sqrt{k}} \left( \frac{1}{\sqrt{k} - u} + \frac{1}{\sqrt{k} + u} \right) \quad (\text{D.2.2})$$

Also note,

$$\frac{1}{u(\sqrt{k} - u)} = \frac{1}{\sqrt{k}} \left( \frac{1}{u} + \frac{1}{\sqrt{k} - u} \right) \quad (\text{D.2.3})$$

and

$$\frac{1}{u(\sqrt{k} + u)} = \frac{1}{\sqrt{k}} \left( \frac{1}{u} - \frac{1}{\sqrt{k} + u} \right) \quad (\text{D.2.4})$$

Putting together the general concepts in Equations ??, ??, and ?? we get...

$$\frac{1}{u(k - u^2)} = \frac{1}{2k} \left[ \left( \frac{1}{u} + \frac{1}{\sqrt{k} - u} \right) + \left( \frac{1}{u} - \frac{1}{\sqrt{k} + u} \right) \right]$$

So, to solve our non-linear differential equation, we make the substitution...

$$\begin{aligned} \int \frac{1}{2k} \left[ \left( \frac{1}{u} + \frac{1}{\sqrt{k} - u} \right) + \left( \frac{1}{u} - \frac{1}{\sqrt{k} + u} \right) \right] du &= \int dt \\ \frac{1}{2k} \int \left[ \left( \frac{1}{u} + \frac{1}{\sqrt{k} - u} \right) + \left( \frac{1}{u} - \frac{1}{\sqrt{k} + u} \right) \right] du &= t + C \\ \frac{1}{2k} (\log |u| - \log |\sqrt{k} - u| + \log |u| - \log |\sqrt{k} + u|) &= t + C \\ \log \left( \frac{|u|}{|\sqrt{k} - u|} \right) + \log \left( \frac{|u|}{|\sqrt{k} + u|} \right) &= 2tk + C \\ \log \left( \frac{u^2}{(|\sqrt{k} - u|)(|\sqrt{k} + u|)} \right) &= 2tk + C \\ \log \left( \frac{u^2}{(|k - u^2|)} \right) &= 2tk + C \\ \frac{u^2}{|k - u^2|} &= e^{2tk+C} = e^{2tk} e^C = C e^{2tk} \\ \frac{|k - u^2|}{u^2} &= C e^{-2tk} \\ \frac{k}{u^2} - 1 &= C e^{-2tk} \\ \frac{k}{u^2} &= 1 + C e^{-2tk} \\ u^2 &= \frac{k}{1 + C e^{-2tk}} \\ u &= \sqrt{\frac{k}{1 + C e^{-2tk}}} \end{aligned}$$

### D.3 Vector and Matrix Calculus

Let  $\mathbf{x}$  be an  $n \times 1$  vector and let  $\mathbf{A}$  be an  $m \times n$  matrix of elements that are not functions of  $\mathbf{x}$ . Then

$$\frac{\partial \mathbf{A}\mathbf{x}}{\partial \mathbf{x}} = \mathbf{A}'. \quad (\text{D.3.1})$$

To see this, we'll just multiply out  $\mathbf{A}\mathbf{x}$  and then apply the definition of vector differentiation (Equation ??). Thus

$$\frac{\partial \mathbf{A}\mathbf{x}}{\partial \mathbf{x}} = \frac{\partial}{\partial \mathbf{x}} \begin{bmatrix} a_{1,1}x_1 + a_{1,2}x_2 + \cdots + a_{1,n}x_n \\ a_{2,1}x_1 + a_{2,2}x_2 + \cdots + a_{2,n}x_n \\ \vdots \\ a_{m,1}x_1 + a_{m,2}x_2 + \cdots + a_{m,n}x_n \end{bmatrix},$$

and, in order to keep our matrix from getting too big for the page, let

$$\begin{aligned} a_{1,1}x_1 + a_{1,2}x_2 + \cdots + a_{1,n}x_n &= \mathbf{a}_1\mathbf{x} \\ a_{2,1}x_1 + a_{2,2}x_2 + \cdots + a_{2,n}x_n &= \mathbf{a}_2\mathbf{x} \\ &\vdots \\ a_{m,1}x_1 + a_{m,2}x_2 + \cdots + a_{m,n}x_n &= \mathbf{a}_m\mathbf{x}. \end{aligned}$$

Thus,

$$\begin{aligned} \frac{\partial}{\partial \mathbf{x}} \begin{bmatrix} a_{1,1}x_1 + a_{1,2}x_2 + \cdots + a_{1,n}x_n \\ a_{2,1}x_1 + a_{2,2}x_2 + \cdots + a_{2,n}x_n \\ \vdots \\ a_{m,1}x_1 + a_{m,2}x_2 + \cdots + a_{m,n}x_n \end{bmatrix} &= \frac{\partial}{\partial \mathbf{x}} \begin{bmatrix} \mathbf{a}_1\mathbf{x} \\ \mathbf{a}_2\mathbf{x} \\ \vdots \\ \mathbf{a}_m\mathbf{x} \end{bmatrix} = \begin{bmatrix} \frac{\partial}{\partial x_1}\mathbf{a}_1\mathbf{x} & \frac{\partial}{\partial x_1}\mathbf{a}_2\mathbf{x} & \cdots & \frac{\partial}{\partial x_1}\mathbf{a}_m\mathbf{x} \\ \frac{\partial}{\partial x_2}\mathbf{a}_1\mathbf{x} & \frac{\partial}{\partial x_2}\mathbf{a}_2\mathbf{x} & \cdots & \frac{\partial}{\partial x_2}\mathbf{a}_m\mathbf{x} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial}{\partial x_n}\mathbf{a}_1\mathbf{x} & \frac{\partial}{\partial x_n}\mathbf{a}_2\mathbf{x} & \cdots & \frac{\partial}{\partial x_n}\mathbf{a}_m\mathbf{x} \end{bmatrix} \\ &= \begin{bmatrix} a_{1,1} & a_{2,1} & \cdots & a_{m,1} \\ a_{1,2} & a_{2,2} & \cdots & a_{m,2} \\ \vdots & \vdots & \ddots & \vdots \\ a_{1,n} & a_{2,n} & \cdots & a_{m,n} \end{bmatrix} = \mathbf{A}' \end{aligned}$$

### D.4

If we let  $\mathbf{x}$  an  $n \times 1$  vector and let  $\mathbf{A}$  be defined in a similar fashion to  $\mathbf{A}$  in Example ??, only now we require  $\mathbf{A}$  to be an  $n \times n$  matrix, then

$$\frac{\partial \mathbf{x}'\mathbf{A}\mathbf{x}}{\partial \mathbf{x}} = (\mathbf{A} + \mathbf{A}')\mathbf{x}, \quad (\text{D.4.1})$$

and if  $\mathbf{A}$  is symmetric, that is  $\mathbf{A} = \mathbf{A}'$ , then  $(\mathbf{A} + \mathbf{A}')\mathbf{x} = 2\mathbf{A}\mathbf{x}$ .

Once again, to understand how these results are derived we will simply multiply out the matrices and then apply the definition of vector differentiation (Equation ??). Thus,

$$\begin{aligned}
\frac{\partial \mathbf{x}' \mathbf{A} \mathbf{x}}{\partial \mathbf{x}} &= \frac{\partial}{\partial \mathbf{x}} \left\{ \begin{bmatrix} x_1 a_{1,1} + \cdots + x_n a_{n,1} & x_1 a_{1,2} + \cdots + x_n a_{n,2} & \cdots & x_1 a_{1,n} + \cdots + x_n a_{n,n} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \right\} \\
&= \frac{\partial}{\partial \mathbf{x}} [x_1(x_1 a_{1,1} + \cdots + x_n a_{n,1}) + x_2(x_1 a_{1,2} + \cdots + x_n a_{n,2}) + \cdots + x_n(x_1 a_{1,n} + \cdots + x_n a_{n,n})] \\
&= \begin{bmatrix} \frac{\partial}{\partial x_1} x_1(x_1 a_{1,1} + \cdots + x_n a_{n,1}) + x_2(x_1 a_{1,2} + \cdots + x_n a_{n,2}) + \cdots + x_n(x_1 a_{1,n} + \cdots + x_n a_{n,n}) \\ \frac{\partial}{\partial x_2} x_1(x_1 a_{1,1} + \cdots + x_n a_{n,1}) + x_2(x_1 a_{1,2} + \cdots + x_n a_{n,2}) + \cdots + x_n(x_1 a_{1,n} + \cdots + x_n a_{n,n}) \\ \vdots \\ \frac{\partial}{\partial x_n} x_1(x_1 a_{1,1} + \cdots + x_n a_{n,1}) + x_2(x_1 a_{1,2} + \cdots + x_n a_{n,2}) + \cdots + x_n(x_1 a_{1,n} + \cdots + x_n a_{n,n}) \end{bmatrix} \\
&= \begin{bmatrix} (2x_1 a_{1,1} + x_2 a_{2,1} + \cdots + x_n a_{n,1}) + (x_2 a_{1,2}) + \cdots + (x_n a_{1,n}) \\ (x_1 a_{2,1}) + (x_1 a_{1,2} + 2x_2 a_{2,2} + \cdots + x_n a_{n,2}) + \cdots + (x_n a_{2,n}) \\ \vdots \\ (x_1 a_{n,1}) + (x_2 a_{n,2}) + \cdots + (x_1 a_{1,n} + x_2 a_{2,n} + \cdots + 2x_n a_{n,n}) \end{bmatrix} \\
&= \begin{bmatrix} (x_1 a_{1,1} + x_2 a_{2,1} + \cdots + x_n a_{n,1}) + (x_1 a_{1,1} + x_2 a_{1,2} + \cdots + x_n a_{1,n}) \\ (x_1 a_{1,2} + x_2 a_{2,2} + \cdots + x_n a_{n,2}) + (x_1 a_{2,1} + x_2 a_{2,2} + \cdots + x_n a_{2,n}) \\ \vdots \\ (x_1 a_{1,n} + x_2 a_{2,n} + \cdots + x_n a_{n,n}) + (x_1 a_{n,1} + x_2 a_{n,2} + \cdots + x_n a_{n,n}) \end{bmatrix} \\
&= \begin{bmatrix} x_1(a_{1,1} + a_{1,1}) + x_2(a_{1,2} + a_{2,1}) + \cdots + x_n(a_{1,n} + a_{n,1}) \\ x_1(a_{2,1} + a_{2,1}) + x_2(a_{2,2} + a_{2,2}) + \cdots + x_n(a_{2,n} + a_{n,2}) \\ \vdots \\ x_1(a_{n,1} + a_{1,n}) + x_2(a_{n,2} + a_{2,n}) + \cdots + x_n(a_{n,n} + a_{n,n}) \end{bmatrix} = (\mathbf{A} + \mathbf{A}') \mathbf{x}
\end{aligned}$$

## D.5

$$\begin{aligned}
&(\mathbf{Y} - \mathbf{X}\hat{\beta})' \left\{ \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} (\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \right\} \\
&= (\mathbf{Y} - \mathbf{X}\hat{\beta})' \left\{ \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \boldsymbol{\theta} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \mathbf{C}\hat{\beta} \right\} \\
&= \mathbf{Y}' \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \boldsymbol{\theta} - \mathbf{Y}' \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \mathbf{C}\hat{\beta} \\
&\quad - \hat{\beta}' \mathbf{X}' \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \boldsymbol{\theta} + \hat{\beta}' \mathbf{X}' \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \mathbf{C}\hat{\beta} \\
&= \hat{\beta}' \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \boldsymbol{\theta} - \hat{\beta}' \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \mathbf{C}\hat{\beta} \\
&\quad - \hat{\beta}' \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \boldsymbol{\theta} + \hat{\beta}' \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \mathbf{C}\hat{\beta} \\
&= 0
\end{aligned}$$

## D.6

$$\begin{aligned}
&[\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} (\boldsymbol{\theta} - \mathbf{C}\hat{\beta})]' [\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} (\boldsymbol{\theta} - \mathbf{C}\hat{\beta})] \\
&= (\boldsymbol{\theta} - \mathbf{C}\hat{\beta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}' \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} (\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \\
&= (\boldsymbol{\theta} - \mathbf{C}\hat{\beta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} \mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} (\boldsymbol{\theta} - \mathbf{C}\hat{\beta}) \\
&= (\boldsymbol{\theta} - \mathbf{C}\hat{\beta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{C}]^{-1} (\boldsymbol{\theta} - \mathbf{C}\hat{\beta})
\end{aligned}$$

## D.7 Proof of $F$ -Distribution

The goal here is to show that Equation ??, that is

$$\frac{(\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta})' [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1} (\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta})/t}{(\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})'(\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})/(n-p)},$$

has an  $F$ -distribution. We will do this by showing that the numerator and the denominator are both  $\sigma^2$  times chi-square variables divided by their degrees of freedom.

Since  $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$  is a linear function of  $\mathbf{Y}$ , and  $\mathbf{Y} \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2)$ , a vector of  $n$  iid random variables, it follows from Equations ?? and ?? that  $\hat{\boldsymbol{\beta}}$  is a vector of  $p$  random variables with a  $N(\boldsymbol{\beta}, \sigma^2[\mathbf{X}'\mathbf{X}]^{-1})$  distribution. Thus, the transformation,  $\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta}$ , results in  $t$  random variables ( $t$  being the number of rows in  $\mathbf{C}$ , the number of tests):

$$\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta} \sim N_t(\mathbf{C}\boldsymbol{\beta} - \boldsymbol{\theta}, \sigma^2\mathbf{C}[\mathbf{X}'\mathbf{X}]^{-1}\mathbf{C}').$$

Under the hypothesis that  $\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta} = 0$  we have,

$$\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta} \sim N_t(0, \sigma^2\mathbf{C}[\mathbf{X}'\mathbf{X}]^{-1}\mathbf{C}'),$$

thus,

$$\begin{aligned} [\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}']^{-1/2}(\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta}) &\sim N_t(0, \sigma^2) \\ \sigma^2[\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}']^{-1/2}(\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta}) &\sim N_t(0, 1). \end{aligned} \quad (\text{D.7.1})$$

Since the sum of  $t$  squared iid  $N(0, 1)$  variables results in a random variable distributed by  $\chi_t^2$ , it follows from Equation ?? that,

$$(\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta})'[\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}']^{-1}(\mathbf{C}\hat{\boldsymbol{\beta}} - \boldsymbol{\theta}) \sim \sigma^2\chi_t^2.$$

Thus, we have shown that the numerator in Equation ?? is  $\sigma^2$  times a chi-square random variable divided by its degrees of freedom.

Showing the same thing for the denominator is a little more tricky as it involves some obscure transformations and knowing a few properties of quadratic forms. Instead of trying to explain the details about quadratic forms that would be required for a full proof, we'll simply go as far as we can with what we have and appeal to your sense of intuition.

Let  $\hat{\boldsymbol{\epsilon}} = \mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}$ , an approximation of the error vector,  $\hat{\boldsymbol{\epsilon}}$ , thus,

$$\begin{aligned} \hat{\boldsymbol{\epsilon}} &= \mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}} \\ &= \mathbf{Y} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} \\ &= (\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\mathbf{Y} \\ &= (\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')(\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}) \\ &= (\mathbf{X}\boldsymbol{\beta} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}\boldsymbol{\beta}) + (\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon} \\ &= (\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon}, \end{aligned}$$

and

$$\begin{aligned} (\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}})'(\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) &= [(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon}]'[(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon}] \\ &= \boldsymbol{\epsilon}'(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')'(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon} \end{aligned} \quad (\text{D.7.2})$$

$$= \boldsymbol{\epsilon}'(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' + \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon} \quad (\text{D.7.3})$$

$$= \boldsymbol{\epsilon}'(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' + \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon} \quad (\text{D.7.4})$$

$$= \boldsymbol{\epsilon}'(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')\boldsymbol{\epsilon}. \quad (\text{D.7.5})$$

Equations ??, ??, ?? and ?? make it clear that  $(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')$  is idempotent, that is,

$$(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')'(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}') = (\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}').$$

We can determine the rank<sup>1</sup> of  $(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}')$  from its trace (that is, the sum the elements on the diagonal) since it is idempotent and symmetric. Using a well known property of traces, that is,  $\text{tr}(ABC) = \text{tr}(CAB)$ , and the fact that  $\mathbf{X}$  is an  $n \times p$  matrix, we have

$$\begin{aligned} \text{tr}(\mathbf{I}_n - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}') &= \text{tr}(\mathbf{I}_n) - \text{tr}(\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}') \\ &= n - \text{tr}(\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}) \\ &= n - \text{tr}(\mathbf{I}_p) \\ &= n - p. \end{aligned}$$

Since  $\epsilon \sim N_n(0, \sigma^2)$ , it follows that  $\epsilon/\sigma \sim N_n(0, 1)$ , thus we can imagine that Equation ?? is the sum of  $n - p$  independent normal random variables. Thus, if we let  $\lambda \sim N(0, 1)$ , then

$$\begin{aligned} \sum_{i=1}^{n-p} \epsilon_i \epsilon_i &= \sum_{i=1}^{n-p} \sigma \lambda \sigma \lambda \\ &= \sum_{i=1}^{n-p} \sigma^2 \lambda^2 \sim \sigma^2 \chi_{n-p}^2 \end{aligned}$$

Showing that the numerator is independent from the denominator also requires some obscure transformations and requires another result from quadratic forms. Without proof, I will state that the following theorem.

Let  $\mathbf{Z}$  is a vector of normally distributed random variables with a common variance and let  $q_1 = \mathbf{Z}'\mathbf{A}\mathbf{Z}$  and  $q_2 = \mathbf{Z}'\mathbf{B}\mathbf{Z}$ , where  $\mathbf{A}$  and  $\mathbf{B}$  are both  $n \times n$  symmetric matrices.  $q_1$  and  $q_2$  are independently distributed if and only if  $\mathbf{A}\mathbf{B} = 0$ .

Now, to show independence, under the hypothesis that  $\mathbf{C}\beta = \theta$ , can re-write the ends of the numerator with

$$\begin{aligned} (\mathbf{C}\hat{\beta} - \theta) &= \mathbf{C}\hat{\beta} - \mathbf{C}\beta \\ &= \mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}[\mathbf{X}'\mathbf{Y} - \mathbf{X}'\mathbf{X}\beta] \\ &= \mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'[\mathbf{Y} - \mathbf{X}\beta] \\ &= \mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\epsilon. \end{aligned}$$

If we let  $\mathbf{A} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}[\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{C}]^{-1}\mathbf{C}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ , then we can rewrite the numerator of Equation ?? as  $\epsilon'\mathbf{A}\epsilon$ . If we let  $\mathbf{B} = \mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ , then we can rewrite the denominator to be  $\epsilon'\mathbf{B}\epsilon$ . Since

$$\begin{aligned} \mathbf{X}'\mathbf{B} &= \mathbf{X}'(\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}') \\ &= \mathbf{X}' - \mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' \\ &= \mathbf{X}' - \mathbf{X}' \\ &= 0, \end{aligned}$$

$\mathbf{A}\mathbf{B} = 0$ , and thus, under the hypothesis, the numerator and denominator are independent.

---

<sup>1</sup>The rank of a matrix is the number of linearly independent rows or columns. For any matrix, the number of linearly independent rows is equal to the number of linearly independent columns.



## D.8

$$\begin{aligned}
 \mathbf{C}\beta\mathbf{U} &= \begin{bmatrix} 1 & 1 & 1 & -1 & -1 & -1 & 0 \end{bmatrix} \begin{bmatrix} \beta_{1,1} & \beta_{1,2} \\ \beta_{2,1} & \beta_{2,2} \\ \beta_{3,1} & \beta_{3,2} \\ \beta_{4,1} & \beta_{4,2} \\ \beta_{5,1} & \beta_{5,2} \\ \beta_{6,1} & \beta_{6,2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\
 &= \begin{bmatrix} (\beta_{1,1} + \beta_{2,1} + \beta_{3,1}) - (\beta_{4,1} + \beta_{5,1} + \beta_{6,1}) & (\beta_{1,2} + \beta_{2,2} + \beta_{3,2}) - (\beta_{4,2} + \beta_{5,2} + \beta_{6,2}) \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\
 &= \begin{bmatrix} (\beta_{1,1} + \beta_{2,1} + \beta_{3,1}) - (\beta_{4,1} + \beta_{5,1} + \beta_{6,1}) \\ (\beta_{1,2} + \beta_{2,2} + \beta_{3,2}) - (\beta_{4,2} + \beta_{5,2} + \beta_{6,2}) \end{bmatrix}
 \end{aligned}$$

## Appendix E

# Alternative Design Matrices for ANOVA

In most text book discussions of design matrices for ANOVA, they commonly dwell solely on what is called the *over parameterized model* and methods for overcoming its limitations instead of the model given in Examples ?? and ?. This is due primarily to the historical origins of ANOVA and reverence to the simplifications that assisted solving the computations by hand. Since we have absolutely no interest in working these problems out by hand, we have adopted a more modern, and in our opinion, more explicit design matrix for ANOVA. However, since it is impossible to avoid these antiquated alternative design matrices and their methods of use, we will describe them here.

Given the data in mice treatment data in Table ??, the over parameterized model is,

$$y = \beta_0 x_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_4 x_4,$$

where  $\beta_0$  estimates a mean value for all of the data (regardless of the particular treatment for each sample) and the remaining four parameters estimate the means of the residuals for each treatment. This model represents the fact that ANOVA was developed prior to the convenient access to computers that we have today and it allowed for the calculations to be done by hand without having to use matrix algebra explicitly. These methods, however, lack generality and obscure the question that you want ANOVA to answer. For example, with the mice data, we ask the question “Are all the treatments the same?” With our modern model, we can easily convert this question into one we can test by asking, “Are the means for each treatment the same?” Using the over parameterized model we end up asking the mildly cryptic question, “Is the variation in the sample due to variation within treatments or variation between treatments?” Both questions eventually will yield the same answer: You decide which one will be easier to explain to someone not already steeped in statistical terminology.

If we are going to use our general hypothesis test (Equation ??) to answer our question with the over parameterized model, we must first create the design matrix. Thus, without displaying the redundant rows, we have:

$$\mathbf{X} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The problem with this design matrix, however, is that it can not be used with our hypothesis testing formula due to the fact that  $\mathbf{X}'\mathbf{X}$  is singular. To work around this problem, statisticians have come up with three solutions. The first is to remove the last column in  $\mathbf{X}$  and modify the parameter vector,  $\boldsymbol{\beta}$  to make up for this change, the second is to modify  $\mathbf{X}$  using what is called  $\sigma$ -restricted notation, and the third is to create a generalized inverse of  $\mathbf{X}$ . Here we will focus on the first two methods since they are encountered most often (see Steel, Torrie and Dickey, for examples using the over parameterized model).

Using the first method we have to make the following changes to the design matrix and the parameter vector:

$$\mathbf{X} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix} \text{ and } \boldsymbol{\beta} = \begin{bmatrix} \beta_0 + \beta_4 \\ \beta_1 - \beta_4 \\ \beta_2 - \beta_4 \\ \beta_3 - \beta_4 \end{bmatrix}.$$

Now, if we multiply  $\mathbf{X}$  and  $\boldsymbol{\beta}$  together, we get:

$$\mathbf{X}\boldsymbol{\beta} = \begin{bmatrix} \beta_0 + \beta_1 \\ \beta_0 + \beta_2 \\ \beta_0 + \beta_3 \\ \beta_0 + \beta_4 \end{bmatrix}$$

Notice that  $\beta_0 + \beta_1$  is just the mean of the first treatment,  $\beta_0 + \beta_2$  is the mean of the second treatment, and so on. Thus, after a lot of work modifying  $\mathbf{X}$  and  $\boldsymbol{\beta}$ , we are exactly where our modern model began.

Using  $\sigma$ -restricted notation, you allow the independent variables to take on three different values, 1, 0 and -1, instead of the binary 1 and 0 used in the other methods. By doing so, we can indicate membership in the last treatment by using -1 for the other treatments. This is because we are assuming that the estimates are unbiased making the sum of the deviations zero. Thus, any particular deviation can be derived from the others as the negative of the sum of the remaining deviations. Our design matrix becomes:


$$\mathbf{X} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & -1 & -1 & -1 \end{bmatrix} \text{ and } \boldsymbol{\beta} = \begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix}.$$

The results here are similar to the over parameterized model.


It is clear that using a *classic* ANOVA approach both obscures the question you are interested in answering and requires more effort on behalf of the individual willing to abide by it. These problems also carry over to ANCOVA whereas our modern model generalizes without any additional effort (see Example ??). Furthermore, since there it is unnecessary to reparameterize the design matrices involved in ANOVA and ANCOVA, we can establish the guideline that any design matrix that *requires* reparameterization should be a signal that you may be making unrealistic assumptions about the nature of the data (see Examples ??, ?? and ??). Thus, the authors are inclined to recommend using our modern approach to ANOVA.

# Appendix F


## 2x2 Factorial Interaction Plots and Their Interpretation




Main Effect Factor A: No  
Main Effect Factor B: No




Main Effect Factor A: Yes  
Main Effect Factor B: No  
Interactions: No




Main Effect Factor A: No  
 Main Effect Factor B: Yes  
 Interactions: No




Main Effect Factor A: Yes  
 Main Effect Factor B: Yes  
 Interactions: No




Main Effect Factor A: Yes  
 Main Effect Factor B: Yes  
 Interactions: Yes



Main Effect Factor A: No



Main Effect Factor A: Yes



Main Effect Factor A: No  
 Main Effect Factor B: No  
 Interactions: Yes

# Appendix G

## License

Copyright (c) 2002 C. Frank Starmer and Joshua D. Starmer. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation;

# Bibliography

- [1] Sen, Ashish and M. Shrivastava (1990). *Regression Analysis: Theory, Methods, and Applications*. New York, NY: Springer-Verlag.
- [2] Steel, Robert G. D., J. H. Torrie and D. A. Dickey (1997). *Principles and Procedures of Statistics: A Biometrical Approach*. Boston, MA: McGraw-Hill.
- [3] Berger, Roger L. (2001). *Experimental Statistics for Biological Sciences*. Raleigh, NC: Self
- [4] Stewart, James. (1999). *Calculus: Early Transcendentals*. Pacific Grove, CA: Brooks/Cole
- [5] Starmer, C. Frank and James E. Grizzle (1968). *A Computer Program for Analysis of Data by General Linear Models*. Chapel Hill, NC: Institute of Statistics Mimeo Series No. 560
- [6] Grizzle, James E., C. Frank Starmer and Gary G. Koch (1969). *Analysis of Categorical Data by Linear Models*. *Biometrics* 15, 489-504
- [7] Leon, Steven J. (2002). *Linear Algebra With Applications*. New Jersey: Prentice Hall
- [8] Casella, George., and Roger Berger. L. (2001). *Statistical Inference: Second Edition*. Pacific Grove, CA: Duxbury
- [9] Nagle, R. Kent, Edward B. Saff and Arthur David Snider. (2000). *Differential Equations and Boundry Value Problems: Third Edition*. Reading, MA: Addison-Wesley.
- [10] Neter et al. (2000). *Applied Linear Statistical Models: Fourth Edition* Boston, MA: McGraw-Hill.
- [11] Kandel, E, Schwartz, J.H. and Jessell, T.M. *Principles of Neural Science. Forth Edition* New York: McGraw-Hill.
- [12] Antonov, I., Antonova, I., Kandel, E.R., and Hawkins, R.D. (2001) *The contribution of activity-dependent synaptic plasticity to classical conditioning in Aplysia* *J. Neuroscience* 21:6413-6422.
- [13] Mayford, M. and Kandel, E.R. (1999) *Genetic approaches to memory storage* *Trends in Genetics* 15:463-470.
- [14] Starmer, C.F. and Grant. A.O. (1985) *Phasic ion channel blockade: A kinetic model and method for parameter estimation* *Molecular Pharmacology* 28:348-356.
- [15] Starmer, C.F. (1988). *Characterizing activity-dependent processes with a piecewise exponential model* *Biometirics* 44:549-559.
- [16] Starmer, C.F., Colatsky, T.J. and Grant, A.O. (2003) *What happens when cardiac Na channels lose their function? 1 - Numerical studies of the vulnerable period in tissue expressing mutant channels.* *Cardiovascular Research* 57:82-91.
- [17] Campbell, N. A., Reece, J. B. and Mitchell, L. G., (1999). *Biology. Fifth Edition* Menlo Park, CA: Benjamin/Cummings



# Index

## Main Ideas

- What is Problem Solving, 4
- Block Designs, 43
- Crafting approximations, 15
- How many experiment must I do?, 6
- Hypothesis Testing, 36
- Reparameterization, 41
- Statistics: using our models, 5
- The essential nonlinearity in biology, 5
- What is a model?, 5
- What is research?, 5
- Why do linear models work?, 5
- Why mess with a Taylor Series, 7