An Integrated, Quantitative Introduction to the Natural Sciences, Part 1: Dynamical Models[1]

---

[1] Notes from the Integrated Science course for first year undergraduate students at Princeton University. The course is an alternative to the introductory physics and chemistry courses, and is intended to encourage students to think quantitatively about a much broader set of phenomena (including, especially, the phenomena of life) than in the usual introductory examples. We assume that students are comfortable with calculus, and have had some exposure to the ideas and vocabulary of high school physics and chemistry. This module about dynamical models comes in the first half of the Fall semester. All inquires can be addressed to `integratedscience@princeton.edu`; current students should refer to `https://blackboard.princeton.edu` for up–to–date course materials.

To the students

Last updated September 1, 2008

As we hope becomes clear, our point of view in this course is rather different from that expressed in conventional introductory science courses. One consequence of this is that we can't simply send you to a standard textbook. These lecture notes, then, are meant something that approximates a book, or more precisely a set of books, plus an extra volume for the labs. You'll see that even the first volume of this project is far from finished. We hope that, as with the rest of the course, you'll view this as a collaborative effort between students and faculty, and give us feedback on what is missing from the notes, or on what needs to be improved.

Relation to the lectures

These notes are not an exact record of the lectures, not least because we hope the lectures (and the lecturers) are still alive enough to be evolving from year to year. We do try to cover the same topics, though, in more or less the same order. Because the match between lecture notes and lectures is loose, we suspect that these notes are not a substitute for the notes which you would take during the lectures. On the other hand, knowing that some of the details are written down here means that you don't have to worry quite so much about writing down every word or equation.

The rough match between lectures and notes, for this first part of the course, including some material to be covered primarily in precepts, is as follows:

0. Introduction
  0.1. A physicist's point of view .... F 12 Sep
  0.2. A biologist's point of view .... M 15 Sep
  0.3. A chemist's point of view .... W 17 Sep

1. Newton's laws, chemical kinetics, ...
1.1 Starting with $F = ma$ .... F 19 Sep
1.2 Boxes and arrows to differential equations .... M 22 Sep
1.3 Radioactive decay and the age of the solar system ..... W 24 Sep
1.4 Using computers to solve differential equations .... (precept)
1.5 Simple circuits and population dynamics .... (precept)
1.6 The complexity of DNA sequences .... F 26 Sep

2. Resonance and response
2.1 The simple harmonic oscillator .... M 29 Sep
2.2 Magic with complex exponentials .... W 1 Oct
2.3 Damping, phases and all that .... F 3 Oct
2.4 Linearization and stability .... M 6 & W 8 Oct
2.5 Stability in a real genetic circuit .... F 10 Oct
2.6 The driven oscillator .... M 13 Oct
2.7 Resonance in the cell membrane .... 15 Oct

3. We are not the center of the universe
3.1 Conservation of energy .... M 20 Oct
3.2 Conservation of $\vec{P}$ and $\vec{L}$ .... W 22 Oct
3.3 Symmetries and conservation laws .... F 24 Oct
3.4 Universality of gravitation .... M 8 Dec
3.5 Kepler's laws .... Tu 9 & W 10 Dec
3.6 Biological counterpoint .... Th 11 Dec

Problems

We cannot possibly overemphasize the importance of derivation as opposed to memorization. Science is *not* a long list of facts, but rather a structure for relating many different observations to one another (more about this soon). Correspondingly, it is not enough for you to learn to recite things that we show you in the lectures; we want you to develop the skills to derive things for yourself, to develop an understanding of how different things are connected. In this spirit, we do *not* want to present a seamless narrative. Rather we want you to pause regularly in the reading of these notes and work things out for yourselves. This is the role of problems, as well as some occasional asides. Thus, instead of collecting the problems at the ends of chapters, as many textbooks, we embed the problems at their appropriate points in the text, encouraging you to stop and think, pick up a pen (or, as needed, put fingers to keyboard), and calculate. Some of the problems are small, essentially asking you to be sure that you understand something which goes by a bit quickly in the text. Others are longer, even open–ended, asking you to explore rather than to find a specific answer.

Students in our course should not be afraid of the large number of problems they find here! We will assign only a fraction of the problems in the weekly problem sets, which will be announced on the `blackboard` web site. The problems which we don't assign may nonetheless prove useful. Short problems could serve as warmup exercises for the assigned problems, while longer problems might serve as fodder for review as exam times approach.

We take this opportunity to remind you that we encourage collaboration among the students in working on the problems, and more generally in learning the material of the course.

Authors

The freshman course evolved through discussions among many faculty. The lectures which form the basis for these notes on dynamics have been given in previous years by William Bialek, David Botstein, John Groves, Michael Hecht, and Joshua Rabinowitz; this year WB, DB & JR are joined by Robert Prod'homme. Much has been added to the presentation by Michael Desai and Matthias Kaschube, who have led precepts. Because the problems play such a central role in the course, we especially thank all of those students who suffered through the early versions, and the succession of teaching assistants who have improved the problems as they prepared solution sets. The text as compiled here is due to WB.

Chapter 0

Introduction

Last updated September 1, 2008

> La filosofia è scritta in questo grandissimo libro che continuamente ci sta aperto innanzi a gli occhi (io dico l'universo), ma non si può intendere se prima non s'impara a intender la lingua, e conoscer i caratteri, ne' quali è scritto. Egli è scritto in lingua matematica, e i caratteri sono triangoli, cerchi, ed altre figure geometriche, senza i quali mezi è impossibile a intenderne umanamente parola; senza questi è un aggirarsi vanamente per un'oscuro laberinto.
>
> G Galilei 1623

Galileo's remarks usually are paraphrased in English as "the book of nature is written in the language of mathematics." Indeed, one way to phrase the goal of physics is that we want to provide a concise and compelling mathematical description of nature, reading and summarizing the grand book to which Galileo refers. Literally the hope is that everything we see around us can be derived from a small set of equations. It is remarkable how much of the world has been "tamed" in this way: We really do know the equations that describe much of what happens around us, and in many cases we actually can derive or predict what we will see by starting with these equations and making fairly rigorous mathematical arguments. This is an extraordinary achievement.

We must admit, however, that much of what fascinates us in our immediate experience of the world remains untamed by mathematics. In particular, life seems much more complex than anything we find in the inanimate world, and correspondingly we suspect that it will be much more difficult to arrive at convincing mathematical theories of biological phenomena. To some extent this suspicion is correct, and our current understanding of biology is much more descriptive and qualitative than is our understanding of the traditional core areas of physics and chemistry. It is important to emphasize, however, that the conventional ways of teaching greatly exaggerate these differences. Thus one teaches physics and chemistry (especially to biologists) by presenting only the very simplest of examples, and one teaches biology by suppressing any role played by quantitative or theoretical analyses in establishing what we actually know. For a variety of reasons, it is time for this to change.

There is a widespread sense that we stand on the threshold of a genuinely new chapter in our understanding of nature, and that our current qualitative descriptions of biological phenomena will be replaced by compelling mathematical theories that are tested and refined through sophisticated quantitative experiments and computational analyses. These extraordinary scientific opportunities require a proportionately radical rethinking of our approach to undergraduate education, and this course is a first effort in this direction. Our goal is to present a more unified view of the sciences as a coherent attempt to discover and codify the orderliness of nature.

It should be clear that, in our view, the current divergence between the teaching of physics and the teaching of biology is a historical artifact which should be remedied. At the same time, the development of separate cultures in the different disciplines is something we should appreciate, respect, and try to understand. But we must always distinguish the current state of the culture from our ultimate goals. This distinction is not universally accepted. A distinguished 20th century biologist (who shall remain nameless here) also quotes Galileo, but he asserts that Galileo was limited to the science of his time, and since he didn't know much biology he of course couldn't realize that biology would be different, and presumably not mathematical. In contrast, we see no reason to doubt Galileo's original assertion that understanding ultimately will mean mathematical understanding.

When we first offered this course in the Fall of 2004, we thought that the goal of unification was so critical that we should stamp out any reference to the historical traditions of different disciplines. This now seems a little naive. Even if we agree that, as a community of scientists, we would like to arrive someday at a seamless understanding of both the animate and inanimate worlds, and that we want this understanding to be faithful to Galileo's vision, *today* the different disciplines are in very different places. As a result, what a biologist sees as the development of a more quantitative biology can be quite different from what a physicist sees as the physics of life, and chemists would have yet a different view. These differences are interesting and important, and the tension among the advocates of the different points of view (perhaps even among your professors in this course!) is a creative tension. Thus, while we want to prepare you for a more unified view of the natural world, we want you to understand the context out which the next generation of scientific developments will come. To do this, we start the course with three lectures that present three different viewpoints: from physics, from chemistry and from biology.[1]

0.1 A physicist's point of view

The search for a compelling mathematical description of nature has led to the invention of new mathematical structures; indeed, starting with calculus, much of advanced mathematics has its origins in efforts to understand nature. One way to organize our exploration is to identify the major classes of mathematical ideas that we use in thinking about the world, and this is the organization that we will follow in this course. But it makes sense to start by outlining the traditional content of a freshman physics course.

Classical mechanics. This is how we predict the trajectory of a ball when we throw it, and how we understand and predict the motions of the planets around the sun; at it's core are Newton's laws. If we broaden our view a bit to include the mechanics of fluids and solids, then this is also the branch of physics that governs the motion of cilia and flagella and much of biological motion on the scale of microns in and around cells; at the other end of the size scale, this is the physics that explains why planes can fly, why bridges support the weight of cars and trucks, and why the weather changes, sometimes unpredictably. In its original form, classical mechanics is *the* great example of describing the dynamics of the world in terms of differential equations, and it is for this purpose that calculus was invented.

Thermodynamics and statistical physics. This part of our subject provides laws of astonishing generality, constraining what can and cannot happen even if we don't understand all the mechanistic details—e.g., the im-

---

[1] In addition, there is the practical fact that you continue your education not as 'integrated scientists' (whatever that might mean), but as physicists, chemists, biologists, ... We want to prepare students with a coherent introduction to the natural sciences that will serve them well no matter which discipline they choose for their major. Thus, this is emphatically *not* just a physics or chemistry course for biology majors, but rather an alternative path into all of the majors. A few brave students have even used the course as a springboard to non–science majors, and one could argue that our integrated view is especially useful for people whose science education will end after their freshman year.

Figure 1: Not the mathematical tools we had in mind when we designed the course.

possibility of perpetual motion machines. This also is the branch of physics that builds a bridge from the image of atoms and molecules whizzing about at random to the orderly phenomena that we see on a macroscopic scale. This is the physics of transistors and liquid crystals, of dramatic phenomena such as superconductivity and superfluidity, and even of the phase transitions that may have driven the extremely rapid expansion of our universe in its initial moments. Here the underlying mathematical structure is probabilistic: What we see in the world are samples drawn at random out of a probability distribution (as when we flip a coin or roll dice) and the theory specifies the form of the distribution.

Electricity and magnetism. It is a remarkable fact that objects seem to interact with each other over long distances, for example two magnets. In the 1800s these interactions were codified by thinking of each object as generating a *field* that pervades the surrounding space, and then the second object responds to the field. Electric and magnetic fields are among the earliest examples of this sort of description, and eventually it was realized that the equations for these fields predict a dynamics of the fields themselves, independent their sources. This dynamics corresponds to propagating waves, much like the classical waves on the water's surface, but the velocity of propagation turned out to be the speed of light, and we now understand that light *is* an electromagnetic wave. Thus electricity and magnetism provides both a great example of how to describe the world in terms of fields and a dramatic example of unification among seemingly disparate phenomena, from the lodestone to the laser.

These three great divisions of our subject thus illustrate three different styles of mathematical description: Dynamical models in terms of differential equations, probabilistic models, and models where the fundamental variables are fields. One of the key ideas that we hope to communicate in this course is that these styles of mathematical description are applicable far beyond their origins. In particular, important parts of chemistry and physics share this underlying mathematical structure, and the same structures are applicable to the more complex phenomena of the living world. Thus instead of organizing our thinking around the historical divisions of physics, chemistry and biology, we will present our understanding of the world as organized by these mathematical ideas.

Dynamical models. Newton's laws predict the trajectories of objects as the solutions to differential equations. Strikingly similar differential equations arise in describing the kinetics of chemical reactions, the growth of bacterial populations, and the dynamics of currents and voltages in electrical circuits. Rather than just teaching these separate subjects, we want to give you an appreciation for the generality and power of these ideas as a framework for understanding dynamical phenomena in nature. We won't stop with the traditional examples, all carefully chosen for their simplicity, but will emphasize that the same approach describes complex networks of biochemical reactions in cells, the rich dynamics of electrical activity in neurons and networks, and so on. In order to meet these goals we will introduce as early as possible the art of approximation and the use of numerical methods as ways of getting both exact answers and better intuition.

Probabilistic models. Freshman physics and freshman chemistry each tackle the conceptually difficult problems of thermodynamics and the statistical description of heat. The mathematical models here are probabilistic—when we make measurements on the world we are drawing samples out of a probability distribution, and the theory specifies the distribution. But Mendelian genetics also is a probabilistic model, and related ideas permeate modern approaches to the analysis of large data sets. Starting with genetics, we will introduce the ideas of probability and proceed through a rigorous view of statistical mechanics as it applies to the gas laws, chemical equilibrium, etc.. Entropy will be followed from its origins in thermodynamics to its statistical interpretation, through its role in information theory and coding, highlighting the startling mathematical unity of these diverse fields. We will introduce the ideas of coarse–graining and approximation, explaining how the same formalism applies to ideal gases and to complex biological molecules.

Fields. While electromagnetism provides a compelling example of field dynamics, the coupled spatial and temporal variations in the concentration of diffusing molecules also generate simple field equations. These equations become richer as we include the possibility of chemical reactions. Surprisingly similar equations describe the migrations of bacterial populations in response to variations in the supply of nutrients, and related ideas describe problems in ecology, epidemiology and even economics (the e–sciences). Even simple versions of the reaction–diffusion problem have important implications for how we think about the emergence of spatial patterns (and, ultimately, body structure) in embryonic development.

Quantum mechanics. The themes of dynamics, fields and probability come together in the quantum world. Rather than a descriptive "modern physics" course, we will present a concise account of the Schrödinger equation, wave functions and energy levels, aiming at a rigorous derivation of atomic orbitals that provides a foundation for discussing the periodic table, the geometry of chemical bonding and chemical reactivity. At the same time we will present some of the compelling paradoxes of quantum mechanics, where we have a unique opportunity to discuss the sometimes startling relationship between mathematical models and experiment.

Reductionism vs. emergence

There is yet another way of organizing our exploration of scientific ideas, and this is the idea of reductionism. Roughly speaking, there is a view of physics and chemistry which says that we start with mechanics of the objects that we see around us, and then we start to take these apart to find out about their constituents. Once we discover that matter is made from molecules, and that molecules are built from atoms, we take apart the atoms to find electrons and the nucleus. Chemistry gives way to atomic physics, and then to nuclear physics as we try understand how the nucleus is built from its constituent parts, the protons and neutrons. Along the way we learn that protons and neutrons are not unique; if they smash into each other with enough energy we see many more particles, each with its own unique set of properties, and indeed many of these exotic particles occur naturally in cosmic rays that rain down on us constantly. Nuclear physics gives way to particle physics, and eventually we understand that the "subatomic zoo" can be ordered by imagining that there are yet more basic constituents called quarks; similarly the electron has several cousins in the particle zoo that are called leptons. The interactions among these particles are described by fields—just as in electricity and magnetism, and indeed there are profound mathematical connections among the equations for the electromagnetic field and the equations for the fields that mediate forces on this subatomic scale. Again the hope is for unification, and to a remarkable extent this has been achieved.

The march from atoms to quarks in one century is one of the great chapters in human intellectual history.[2] At the same time, describing physics *only* as the constant drive to peel away layers of description, searching for the 'fundamental' components of the universe, is a bit simplistic. In some circles this reductionist drive is emulated—trying to make some area of science "more like physics" often seems to mean making it more reductionist. At the same time, for many people "reductionist" is sort of a dirty word: Surely we must recognize that systems are more than the sum of their parts, and that essential functions are lost when we tear them to bits in the process of understanding them.

In fact, the portrait of physics as a reductionist enterprise misses about half of what physicists have been doing since ∼1950. The other half of physics is about how macroscopic or collective properties emerge from the interactions among more elementary constituents. A familiar (if, in the end, somewhat complicated) example is provided by water. We know that pure water is made from only one kind of molecule, and the essence of the reductionist claim is that the properties of water are determined by its molecular composition. But the most obvious properties of liquid water—-it feels wet, and it flows—clearly are not properties of single water molecules. It's not even clear that a small cluster with tens of molecules would have these properties; when we look at small numbers of molecules, for example, our attention is immediately drawn to the fact that there is a lot of empty space in between the molecules, while on a human scale this space is imperceptible and the water seems to be a continuous substance.

The discrepancy between single molecules and macroscopic properties is even greater if we think about solid water (ice). The statement that an ice cube is solid is a statement about rigidity: if we push on one side of the ice cube, the whole cube moves together, and in particular the opposite face

---

[2] As new Princetonians, you can take pride in knowing that some of the important steps in this grand march took place right here.

of the cube also moves. This property can't even be defined for a single molecule, and it certainly isn't something that happens if we just have a few molecules. Indeed, it's not obvious why it works at all. Imagine lifting a one kilogram cube of ice. You grip the edges, apply a force, and the entire block is raised by say one meter. In the process you have put about 10 Joules of energy into work against the force of gravity. This is more than enough energy to rip the first layers of water molecules off of the faces of the block, but that's not what happens. Instead of flying apart in response to the force, or flowing around your fingers like liquid water, all of the water molecules in the block of ice (even the ones that your not touching!) move together. This rigidity or solidity of ice clearly has something to do with the water molecules, but clearly involves the whole block of material being something more than just the "sum of the parts."

There is a whole language for talking about the block of ice or the flow of liquid water (or the flow of air—wind—in the atmosphere) that is different from how we talk about molecules. Thus we can talk about a tornado as if it were an object moving across the globe, or about the hardness of the ice surface, and neither of these things correspond in a simple way to the things we might measure for individual molecules. Similar things happen in other materials, sometimes quite dramatically. Thus, while a single electron moving through a solid might rattle around, bumping into various atoms and eventually losing its way, when we cool a hunk of metal down to very low temperatures, all the electrons can flow together in an electrical current that lasts essentially forever, a phenomenon called superconductivity. If each molecule in a liquid crystal display responded individually to applied electric fields, you could never get the bright and brilliant colors that you see on your laptop computer; again there is some collective behavior in large groups of molecules, and the whole is more than the sum of the parts.

The understanding of how these sorts of macroscopic, collective behaviors emerge from the dynamics of electrons, atoms and molecules did not come in one bold stroke, as in the high school version of science history where Einstein writes down the theory of relativity and suddenly everything changes. Instead there were many independent strands of thought, which started to come together in the 1960s.[3] Gradually it became clear that, at

---

[3] One of the landmarks was a lecture given in 1967 by our own Phil Anderson, published some years later: PW Anderson, More is different. *Science* **177,** 393–396 (1972). The piece was written partly in opposition to the notion that the reductionist search is somehow 'more fundamental' than the search for synthesis, and Phil's unique combative styles comes through clearly in his writing. Sociology aside, this paper gives a beautiful statement of the idea hinted at above, that even our language for describing things evolves as we move

least in a few cases, it was possible not just to understand that collective effects happen, not just to describe these more macroscopic dynamics, but even to understand how they can be derived, at least in outline, from something more microscopic. In this process, something remarkable happened: we understood that when macroscopic collective effects arise, once we focus on these effect we lose track of many of the microscopic details. The flow of a fluid provides a good example, since the equations which describe this flow depend on the density and viscosity of the fluid, and essentially nothing else. All of the complicated properties of the molecules—the geometry of their chemical bonds, their van der Waals interactions and hydrogen bonds with each other, ... —none of this matters except to set the values of those two parameters.[4] The 1970s brought even more dramatic examples of such 'universality,' including building precise mathematical connections between seemingly completely different phenomena, such as the liquid crystals and superconductors mentioned above, but this is going too far for an introduction.

We'll come back explicitly to these ideas in the second half of the Fall semester, and more generally our discussion of statistical physics will be all about how to build up from the microscopic description of atoms and molecules to the phenomena that we observe on a human scale. Perhaps we can even give a hint of how the ideas of universality have given physicists the courage to write down theories for much more complex phenomena, up to the phenomena of perception and memory in the human brain. For now, however, enough philosophy.

The simplest models

All of the models mentioned above involve (at least) some calculus. But there is a much simpler class of model, maybe one that is so simple we don't even think of it is a "model of nature" in the grand sense that we use the

---

from one level (e.g., molecules) to the next (fluids and solids). There can be independent discoveries of the relevant mathematical laws at each level, and it is by no means obvious how to move from one level to the next. Indeed, in those cases where we learn how to derive laws at one level from the dynamics 'underneath,' it is a great triumph.

[4] Actually one can do more. These parameters have units, and nobody tells you what system of units to use. By adjusting your system of units, you can almost make the parameters all be equal to one, so that there is truly nothing left of the molecular details. This isn't quite right, because there is still a dimensionless ratio of parameters that combines the properties of the fluid with the typical spatial scale and speed of the fluid flow, but this one number (Reynolds' number) tells the whole story. Flows with the same Reynolds' number look *quantitatively* the same no matter what molecules make up the fluid.

word today: models in which one variable is just a linear function of the other. There are many familiar examples:[5]

- The voltage drop V across a resistor is proportional to the electrical current I that flows through the resistor, $V = IR$. This is Ohm's law, and R is the resistance.
- The force F that we feel when we stretch a spring is proportional to the distance x that we stretch it, $F = -\kappa x$. This is Hooke's law, and κ is called the stiffness of the spring.[6]
- The charge Q on a capacitor (more precisely, the charge difference between the plates of the capacitor) is proportional to the voltage difference V, $Q = CV$, where C is the capacitance.
- When an object moves through a fluid at velocity v, it experiences a drag force $F = -\gamma v$, where γ is called the drag or damping constant.
- The force of gravity is proportional to the mass of an object, $F = mg$ (this one is special!).

There are even more examples, such as the fact that the rate at which your coffee cools is proportional to the temperature difference between the coffee and the surrounding air. While simple, there really is a lot going on in these "laws."

First of all, the notion that these are "laws" needs some revisiting. Obviously if you take a rubber band and pull, then pull some more, eventually the rubber band will snap (don't blame me for the bruise if you feel compelled to verify this). Certainly when the band breaks Hooke's law stops being valid, but we suspect that the force stops being proportional to the distance we have stretched long before the band breaks. These considerations describe what happens when we stretch the rubber band, but of course we can also compress it, and then we know that it can go slack, so that no force is required to bring the ends together. See Fig 3. So in what sense is Hooke's law a law?

Statements such as Hooke's law and Ohm's law are very good approximations to the properties of many real materials, but they are not "laws"

---

[5] Hopefully these are reminders of things you've learned in your high school course. If you miss one, don't worry, most will come back later in more detail.

[6] The symbol κ is the Greek letter "kappa," and (below) γ is "gamma," both lower case. It's incredibly useful to know the Greek alphabet; it's not hard, and if you find yourself in Athens you'll be able to read the street signs. See Fig 2.

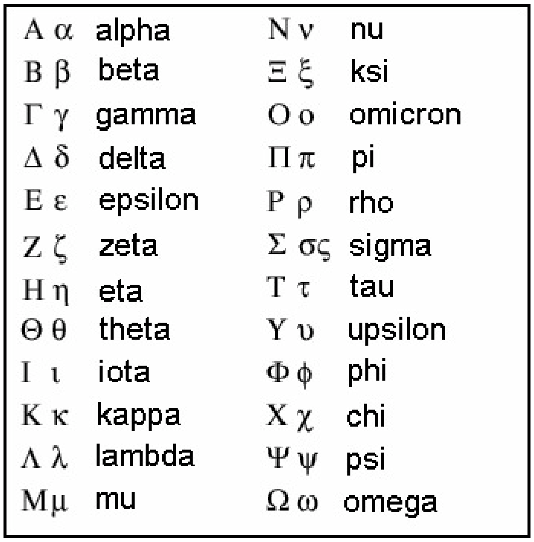

Α α	alpha	Ν ν	nu
Β β	beta	Ξ ξ	ksi
Γ γ	gamma	Ο ο	omicron
Δ δ	delta	Π π	pi
Ε ε	epsilon	Ρ ρ	rho
Ζ ζ	zeta	Σ σς	sigma
Η η	eta	Τ τ	tau
Θ θ	theta	Υ υ	upsilon
Ι ι	iota	Φ φ	phi
Κ κ	kappa	Χ χ	chi
Λ λ	lambda	Ψ ψ	psi
Μ μ	mu	Ω ω	omega

Figure 2: The Greek alphabet: upper case, lower case, and the name of the letter. From `http://gogreece.about.com/`. We'll usually write the lower case sigma as σ.

of universal applicability such as Newton's $F = ma$. But why do these approximations work? Let's think about some function, perhaps force as a function of length for the spring, or voltage as a function of current in a wire (resistor), etc.. Let's call this function $f(x)$. If you can see the whole function, it might be quite funny looking, as in Fig 4. On the other hand, if we only want to know the value of the function close to the place where $x = x_0$, we can make approximations (Fig 5).

We could start, for example, by ignoring the variations all together and saying that since x is close to x_0, we'll just pretend that the function is constant and equal to $f(x_0)$. A bit silly, perhaps, but maybe not so bad. The next best thing is to notice that you can fit a straight line to the function in the neighborhood of x_0, which is the same as writing

$$f(x) \approx f(x_0) + a(x - x_0), \tag{1}$$

where a is some constant that measures the slope of the line.

The next approximation would be to fit a curve in the neighborhood of x_0, starting with a parabola,

$$f(x) \approx f(x_0) + a(x - x_0) + b(x - x_0)^2. \tag{2}$$

Clearly we could keep going, using higher and higher order polynomials to try and describe what is going on in the graph. Notice that by writing our approximations in this way we guarantee that they give exactly the right answer at the point $x = x_0$.

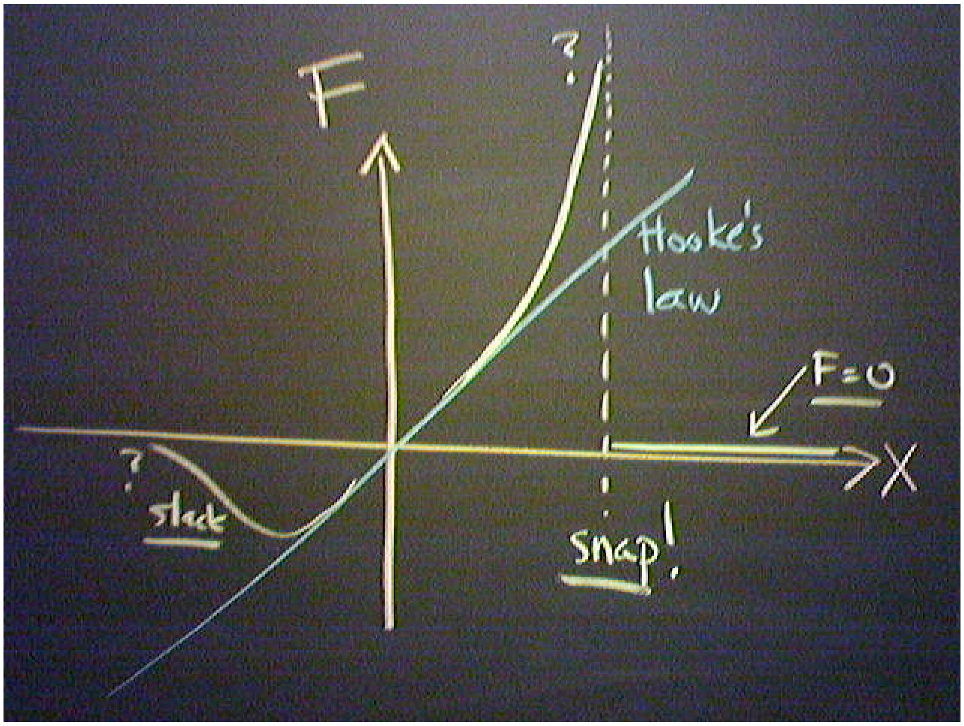

Figure 3: Force F that *opposes* lengthening of a rubber band, versus the length x of the band. Straight blue line indicates Hooke's law, $F = \kappa x$. At some critical length the band will snap; leading up to this it takes extra force, although it's not obvious exactly what happens, but after the snap it takes no force to move the ends apart since they're not connected (!). At the other side, once the rubber band shortens to the point of going slack, the force again goes to zero.

Some of you will recognize that what we are doing is using the *Taylor series* expansion of the function $f(x)$ in the neighborhood of $x = x_0$. We recall from calculus that for any function with reasonable smoothness properties, for some range of x in the neighborhood of x_0 we can write

$$\begin{aligned} f(x) &= f(x_0) + \left[\left. \frac{df(x)}{dx} \right|_{x=x_0} \right] \cdot (x - x_0) \\ &\quad + \frac{1}{2} \left[\left. \frac{d^2 f(x)}{dx^2} \right|_{x=x_0} \right] \cdot (x - x_0)^2 \end{aligned}$$

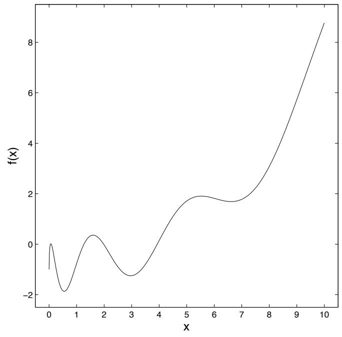

Figure 4: A function $f(x)$, chosen to have some interesting bumps and wiggles.

$$+\frac{1}{3!}\left[\left.\frac{d^3 f(x)}{dx^3}\right|_{x=x_0}\right]\cdot(x-x_0)^3+\cdots, \tag{3}$$

where $\cdots$ are more terms of the same general form; we can write the same equation as

$$f(x)=f(x_0)+\sum_{n=1}^{\infty}\frac{1}{n!}\left[\left.\frac{d^n f(x)}{dx^n}\right|_{x=x_0}\right]\cdot(x-x_0)^n. \tag{4}$$

Now is a good time to be sure that you remember how to read and understand the summation symbol, as well as the vertical bar that means "evaluated at."

Figure 5 makes clear that the Taylor series works in a practical sense: We can start with a pretty wild function, and if we focus our attention on a small neighborhood then the first few terms of the Taylor series are enough to get pretty close to the actual values of the function in this neighborhood. Notice that this practical view is different from what you may have learned in your calculus class, where the emphasis is on proving that the Taylor series converges, i. e. that if we keep enough terms in the series we will eventually get as close as we want. The idea here is that just the first couple of terms are enough as long as we don't let $|x-x_0|$ get to be too big.

We can think of "laws" like Hooke's law or Ohm's law as the first terms in a Taylor series. Thus, as in Fig 3, the real relation between force and length (or voltage and current) might be quite complicated, but as long as we don't pull or push too much, the linear approximation works. But we haven't said what "too much" means. Look back at Fig 4. Although the function has lots of bumps and wiggles, they don't come very close together. Thus, as we sweep out a range of $\Delta x\in[-1,+1]$, things look pretty smooth. In fact Fig 5 focuses on a region of this size, and inside this region a low order approximation indeed works very well. So we understand that the

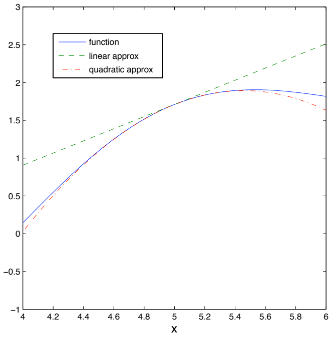

Figure 5: A closer view of the function $f(x)$ from Fig 4 in the neighborhood of $x = 5$, together with linear and quadratic approximations.

first terms of a Taylor series are enough if we look at a range of x that is smaller than some *natural scale* of variations in the function we are trying to approximate.

A key point in physical systems is that the "natural scale" we are looking for has units! When we say that Hooke's law is valid if we don't stretch the spring too much, how far is too much corresponds to some real physical distance. What is this distance? Similarly, when we pass current through a wire, we say that Ohm's law is valid if we don't try to use too much current or apply too large a voltage ... but what is the natural scale of current that corresponds to "too much"? Being able to answer these questions is a critical step in thinking quantitatively about the natural world, and we will return to these problems several times during the course.

These first problems should be thought of as warmup exercises, to get you back in the rhythm of problem solving.

Problem 1: In order to answer questions about the "natural scale" for different phenomena, you will need to think about orders of magnitude. It might be hard to explain why something comes out to be exactly 347 (in some units), but you should be able to understand why it is ~ 300 and not ~ 30 or ~ 3000. Indeed, sometimes it's more satisfying to have a short argument for the approximate answer than a long argument for the exact answer.

(a.) What is the typical distance between molecules in liquid water? You should start with the density of water, $\rho = 1\,\text{gram}/\text{cm}^3$.

(b.) Many bacteria are roughly spherical, with a diameter of $d \sim 1\,\mu\text{m}$. If you divide up the weight of the bacterium, you find that it is 50% water, 30% protein, and 20% other molecules (e.g., RNA, DNA, lipid). A typical protein has a molecular weight of 30,000 atomic mass units (or Daltons).[7] Roughly how many protein molecules make up a bacterium? A typical bacterium has genes that code for about 5,000 different proteins. On average, how many copies of each protein molecule is present in the cell?

---

[7] Recall that one atomic mass unit is a mass of one gram per mole.

(c.) In [b] you computed an average number of copies for all proteins, but different proteins are present at very different abundances inside the cell. Indeed, there are important proteins (such as the transcription factors that help to turn genes on and off) that function at concentrations of $\sim 1 - 10\,\mathrm{nM}$. How many molecules of these proteins are present in the cell?

(d.) To encourage this kind of thinking, Enrico Fermi famously asked "How many piano tuners are there in America?" during a PhD exam in Physics. Similar questions include: How many students enter high school in the United States each year? How many college students each year need to become teachers in order to educate all these people? How many houses does the tooth fairy visit each night?[8] Answer these questions, and formulate one of your own.

Problem 2: If we have a block of material with area A and length L, then the stiffness for stretching or compressing along its length will be $\kappa = YA/L$, where Y is called the Young's modulus.

(a.) Explain why the stiffness should be proportional to area and inversely proportional to length.

(b.) Show that Y has units of an energy density (or energy per unit volume—joules/m^3 or erg/cm^3). Note that this makes sense because the energy that we store in the block when we stretch it by an amount ΔL, $E_{\rm stored} = (1/2)\kappa(\Delta L)^2$, works out to be proportional to the volume ($V = AL$) of the block:

$$E_{\rm stored} = \frac{1}{2}\kappa(\Delta L)^2 = \frac{1}{2}Y \cdot (AL) \cdot \left(\frac{\Delta L}{L}\right)^2. \tag{5}$$

(c.) Diamond is one of the stiffest materials known, and it has $Y \sim 10^{12}\,\mathrm{N/m^2}$ (or $\mathrm{J/m^3}$). The density of diamond is $\rho = 3.52\,\mathrm{g/cm^3}$. Convert Y into an energy per carbon atom in the diamond crystal. How does this compare with the energy of the chemical bonds in the diamond crystal? Note that you'll need to look up this number ... be careful about units! Does your answer make sense?

Problem 3: Going back to the *E coli* in Problem 1, we want to understand the implications of the fact that one bacterium can make a complete copy of itself (dividing into two bacteria) in $\tau \sim 20\,\mathrm{min}$.

(a.) Proteins are synthesized on ribosomes, which can add ~ 20 amino acids per second to a growing protein chain. If the typical protein has 300 amino acids, how many protein molecules can one ribosome make within the doubling time τ?

(b.) In order to double within τ, the bacterium presumably has to make an extra copy of all of its protein molecules. How many ribosomes does it need in order to do this? Make use of your results from Problem 1 on the numbers of protein molecules per cell.

(c.) The ribosome is quite large as molecules go, with a diameter of $\sim 25\,\mathrm{nm}$. If you could cut open the bacterium and see all the ribosomes, how far apart would they be? Is there much empty space between the ribosomes, or is the cell's interior more densely packed?

It would be wrong to leave this discussion of very simple models without talking about one of the very simplest—so simple, in fact, that from our

---

[8] Admittedly, this is a bit more hypothetical than Fermi's problem.

modern perspective we can miss that it is a model at all. You all know that

$$2H_2 + O_2 \rightarrow 2H_2O. \tag{6}$$

We talk about this equation in terms of each water molecule being made out of two hydrogen atoms and one oxygen atom. But not so long ago, we didn't know about atoms. What we did know was that if you mix hydrogen and oxygen together, you get water. When people looked more carefully, they found that this could be made quantitative: A certain amount of hydrogen and oxygen, in certain proportions, are needed to make a certain amount of water. But "amounts" are measured in some units which, from our modern point of view, are rather arbitrary—liters (or worse, gallons) of the gases, grams of the liquids or solids, ... So you might learn that some number of liters of hydrogen gas plus some number of liters of oxygen gas produces some number of grams of liquid water.

$$44.8 \text{ liters } H_2 + 22.4 \text{ liters } O_2 \rightarrow 36 \text{ grams } H_2O. \tag{7}$$

Now you do other experiments, mixing hydrogen and oxygen with other materials. Each time there is some rule about how the different amount combine. It really was an amazing discovery that if you choose your units correctly you can turn all these funny numbers [as in Eq (7)] into the integers of Eq (6). Thus, if you say that one unit of gas at room temperature and pressure is 22.4 liters, then all of the reactions involving gases simplify, and so on.

What's going on here? Now we know, of course, that we should measure the number of molecules or the number of moles of each substance, and then the rule for combining macroscopic quantities just reflect the rules for combining individual atoms. But we can think about trying to write down the rules for the macroscopic quantities alone, and then these are simple functional relations—in fact they are linear relations, not unlike Hooke's law. The wonderful thing is that the coefficients in these linear relations don't take on arbitrary values (as with the stiffness of a spring) but if we choose our units correctly these coefficients are just pure numbers. It took a long time to go from this discovery to the modern view of atoms and molecules, but when you find that there is a way of looking at the world in which the numbers you need to know are integers, you know you're on to something!

What physicists do

`I'd like to add a little more here, to bring the text back to what physicists do today ...`

0.2 A chemist's point of view

0.3 A biologist's point of view

`Josh & David, would you like to put something here in print?`