Part I
In the Beginning: Generating, Detecting, and Manipulating the MR (NMR) Signal

1
Laying the Foundation: Nuclear Magnetism, Spin, and the NMR Phenomenon

The Overall Aim

If we bring together a group of people involved in magnetic resonance imaging (MRI) and present them with a magnetic resonance (MR) image, each will see something different in that image. Radiologists will hone in on subtle pathology, psychiatry researchers will notice minute asymmetries in cortical sulci, technologists may pick up on poor positioning, and physicists will evaluate signal-to-noise and detect artifacts. What are we all trying to achieve at the end of the day? Regardless of the type of image acquired or the purpose for which it was acquired, our single common goal in MRI is this:

Differentiate the tissue in two adjacent locations based on the way that tissue behaves in the MRI environment.

If the signal extracted from those two locations is identical, we cannot differentiate the two tissues from each other. This could be because each location does in fact contain the same tissue (e.g., two locations within the cortex of the kidney) or because the MR image was not made sensitive to the difference between the two tissues. In this latter case, a tumor could be virtually indistinguishable from the normal tissue within which it is growing.

To accomplish our goal of differentiating adjacent tissues of different composition, two tasks must be accomplished in properly designing the MRI examination: (1) resolving the tissues to unique locations in the image and (2) detecting different signal amplitude (intensity) from the two tissues.

Spatial Resolution

The image must be physically capable of resolving the two locations of interest in the tissue to distinct locations within the image; that is, the image must have sufficient *spatial resolution*. An MR image represents a slice of tissue with a defined thickness. This slice is then divided into a two-dimensional array of prism-shaped pieces that we call voxels (for "volume elements"). If we are trying, for example, to differentiate a 5-mm nodule in the liver of a cancer patient from the surrounding normal liver tissue, the voxel must be close to or, preferably, smaller than 5 mm. This is because each voxel is sampled as a single MRI signal that is an average of signal arising from *all* of the tissue within the voxel. In our example liver nodule, if the voxel is so large, say 10 mm, that it contains more normal tissue than abnormal, the average signal we sample will be dominated by normal tissue. Because this signal is no different than that arising from adjacent voxels containing normal liver, the nodule may go undetected.

Contrast Resolution

The MRI acquisition must also elicit a different signal from each of the tissues we wish to separate. That is, the image must have sufficient *contrast resolution*. Even if the image has exceptional spatial resolution (i.e., very small voxels much smaller than the lesion of interest), if both the nodule and normal liver yield the same signal, we will be unable to differentiate the nodule from normal liver tissue. In an effort to maximize the difference in signal between tissues, we modulate the measured signal by adjusting the parameters of the MR acquisition.

As we progress in our understanding of MRI, do not lose sight of the fact that all we are ever trying to do is to differentiate two adjacent tissues based on their different MR signal. The same is true for any and all MRI images whether spin echo, diffusion weighted, or even functional MRI.

Where Does the MRI Signal Come From?

Before we venture into the realm of imaging, we must first understand how the MR signal is generated and measured. To keep things manageable, the following disclaimer should be kept in mind: *Until further notice*, we will limit our discussion to the measurement of signal from a homogeneous sample (maybe a bowl of Jell-O?). *No* attempt will be made, at this point, to determine where in the sample the signal comes from; the signal comes from everywhere in the entire sample.

Nuclear Magnetic Resonance

Nuclear magnetic resonance (NMR) is a physical phenomenon that occurs when *certain* elements interact with a magnetic field. NMR is the process by which the signal detected in MRI is generated; it is the foundation on which MRI is built. Some common elements that demonstrate NMR are

Table 1-1. Some elements that undergo NMR.

^{1}H	^{31}P
^{13}C	^{19}F
^{15}N	^{129}Xe

listed in Table 1-1. In order to qualify for this list (and other elements and isotopes do qualify), the element must have a nonzero magnetic moment. It is not necessary for us to delve into what a nonzero magnetic moment is, but it will be present when either the number of protons *or* neutrons in an atom is odd. Thus, helium can never undergo NMR because its nucleus is composed of two protons and two neutrons. So why is it so important that the atom have an odd number of protons or neutrons? A truthful answer requires a discussion involving quantum mechanics. We will touch on this a bit in the following sections. The bottom line, however, is only those nuclei with nonzero magnetic moments demonstrate a property called *spin angular momentum*. The essential role of spin angular momentum in NMR will be introduced shortly.

Spin Semantics

Although many nuclei can undergo NMR, we will confine our discussion to the hydrogen nucleus (^{1}H). The type of sample we will be imaging and the question that we wish to answer directs the choice of nucleus to be examined. In human tissue, which is composed largely of hydrogen-containing water (H_2O), hydrogen is the most abundant of all the NMR-capable nuclei. For this reason, human MRI is focused almost exclusively on hydrogen. However, if we wished to look at energy metabolism and ATP, ^{31}P would be the nucleus of choice. If we were interested in glucose metabolism, ^{13}C would be best.

Because the hydrogen nucleus contains a single proton and nothing more, hydrogen nuclei are also referred to merely as

What about the electrons?

As small charged subatomic particles, electrons also demonstrate magnetism and spin. Thus, they demonstrate both magnetism and spin angular momentum. Why, then, are we ignoring them in our discussion of NMR? Because of their much smaller mass, electrons have a much higher gyromagnetic ratio and precess at frequencies in the gigahertz range; signals in this frequency range will not be detected by our detection hardware (see Chapter 5). Additionally, their resonance signal—called electron spin resonance, or ESR—is comparatively weak relative to the NMR signal. The magnetic fields expressed by electrons, however, are quite important, as they induce variability in the static magnetic field of the MRI scanner (B_0), causing the chemical shift effects so important in spectroscopy and imaging.

protons. As we will see shortly, the nucleus must have a quantum mechanical property termed *nuclear spin*. For this reason, nuclei are also commonly referred to merely as spins. Note that in the case of NMR of nuclei other than ^{1}H, the term *proton* would be incorrect, whereas the term *spin* would be appropriate.

For human MRI, the terms *nucleus*, *proton*, and *spin* are interchangeable.

Prerequisites to NMR: Nuclear Magnetism

NMR is based on the presence of two properties of the atom: (1) magnetism and (2) spin angular momentum. The protons within the nucleus of any atom contain electric charge and generate a magnetic field. Whereas it is tempting to think of such nuclear magnetism in terms of classic electromagnetic phenomena with spinning charges generating a magnetic field based on Faraday's law of induction, nuclear magnetism is in fact a quantum mechanical phenomenon, and nuclear particles do not, as currently understood, actually spin in the physical sense.

Nuclear magnetism does, nonetheless, result in a very real magnetic field that is local to the nucleus and behaves just like the magnetic field of a permanent magnet or compass needle. Magnetic field lines describe this nuclear magnetic field and are identical to those generated by a common bar-shaped permanent magnet (Fig. 1-1).

Basic Electricity and Magnetism: *Faraday's Law of Induction*

Moving charge (i.e., electricity) will induce (generate) a magnetic field. It really is that simple. The key to understanding NMR, however, is that *Faraday's law also works in reverse! Moving magnetic fields induce electricity—voltage and/or current in a conductor.*

A real-life example: In a hydroelectric dam, the rushing water pushes turbines that spin large magnets nestled within huge coils of wire. The moving magnets induce voltage in these wires, ultimately sending electric current through the power lines that provide the electricity that runs your toaster. *Remember: We require magnets in motion. In the case of the hydroelectric dam, the magnets are permanent; in MRI, they are water protons.*

Field lines are not something magical. They are very real and a useful way to describe the strength, orientation, and homogeneity of a magnetic field (see Chapter 5). For our descriptions of nuclear magnetism, however, it will suffice to describe the magnetic field of the nucleus using a single vector. The orientation of this magnetic field vector will indicate the orientation of the nuclear magnetic field, and the length of the vector will indicate the strength of the field. In the absence of any magnetic field external to the nucleus, orientation of the nuclear magnetic field will be random. In

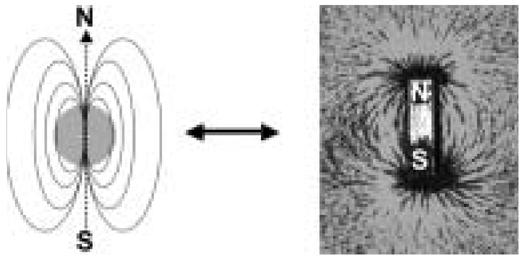

FIGURE 1-1. The proton's magnetic field. The magnetic field lines of a common bar magnet (right) are shown by the distribution of iron filings. The magnetic field of a proton (left) behaves in the same manner, exhibiting field lines and polarity identical to those produced by the bar magnet.

the presence of an externally applied magnetic field, however, the nuclear magnetic field will align with the externally applied magnetic field, much as a compass needle aligns with the earth's magnetic field.

At this point, we have described the magnetism of the ^{1}H nucleus and can see that our proton essentially behaves as a tiny magnet, aligning with an external applied magnetic field. To demonstrate NMR and be useful for MRI, however, the proton must also have *spin angular momentum*.

Prerequisites to NMR: Nuclear Spin Angular Momentum

Spin angular momentum is a property of certain nuclei (those that exhibit NMR). Its existence becomes apparent when we observe the interaction of such a nucleus with an externally applied magnetic field. Though not precisely applicable to the world of small particles like atomic nuclei, a real-world example can at least give us a feeling for what angular momentum is. When a figure skater crouches into a high-velocity spin, the skater may wobble. If we ask the skater what he or she felt during that spin, they will describe a centrifugal force pushing his or her body away from its vertical alignment. In fact, any Olympic hopeful skater will tell us that it takes substantial energy to resist this force and maintain the spin. Similarly, if we take a small weight, tie it to the end of a string, and whirl the string and weight overhead like a lasso, what happens? The string remains taut. Again, this centrifugal force is at work. The physical phenomenon that produces this centrifugal force is called *angular momentum*. The same force causes the wobble of a spinning top.

Spin angular momentum is an analogous force that is applicable to small particles such as atomic nuclei. Consider another macroscopic example. Tie a permanent magnet to a string and suspend it a few centimeters above a tabletop; the entire system will lie at rest with the string plumb to gravity. Next, place a second permanent magnet on the tabletop and slide it under the suspended magnet. What happens? The suspended magnet will begin to swing from side to side over the magnet placed below it on the tabletop (Fig. 1-2A). The magnet below attracts the suspended magnet, causing it to move closer. Momentum of the suspended magnet causes it to overshoot the location of the magnet below. Next, the suspended magnet is drawn back toward the magnet on the tabletop. Ultimately, the suspended magnet swings back and forth in a straight line until it eventually slows and comes to rest over the magnet that was placed on the tabletop. Why does the suspended magnet stop swinging? The system comes to rest only because of friction within the suspensory mechanism. If a completely frictionless system could be developed to suspend the magnet, it would swing back and forth forever! By the way, although the direction in which the magnet swings is completely unrestricted, it will swing back and forth in a straight line.

Let's complicate the system slightly. Install a small motor at the point where the string is suspended and rotate the string at this point. Note that the string is still free to swing back and forth. What happens now? Instead of swinging back and forth in a straight line, the spinning magnet moves in a circular trajectory around the magnet lying on the table below (Fig. 1-2B). If we would map the path of the suspending string, it would delineate a cone. This change from a linear to a circular path is due only to the pres-

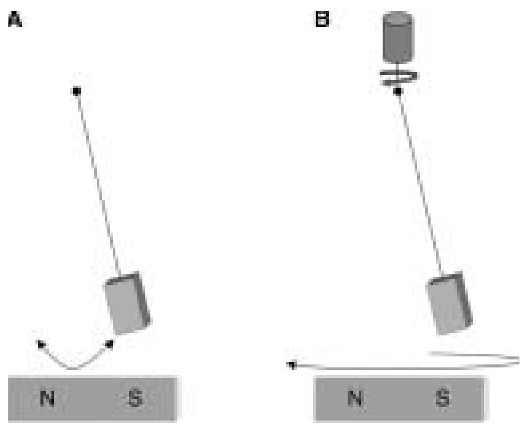

FIGURE 1-2. Spin angular momentum. (A) The magnet suspended from a universal joint is free to swing in any direction. (N and S refer to the poles of the magnet.) When a second magnet is placed underneath it, the suspended magnet will swing back and forth through a plane parallel to the magnetic field of the stationary magnet. Only because of friction at its point of attachment will the suspended magnet eventually come to rest. (B) When the point of attachment is connected to the shaft of a small motor, angular momentum is added to the system. When the second magnet is positioned underneath, the suspended magnet circles so that the string to which it is attached traces the surface of a cone.

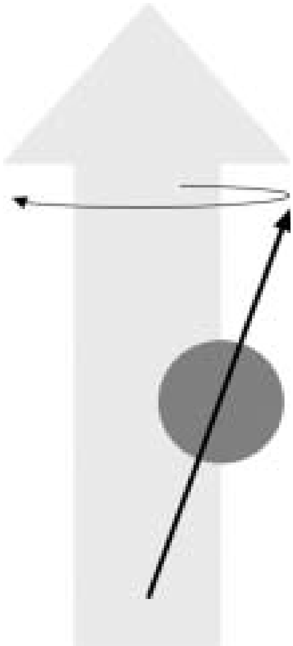

FIGURE 1-3. Precession. Interaction of the magnetic field of the proton with an applied magnetic field leads to precession. The vector describing the proton's magnetic field circles that describing the static magnetic field. Its path traces the surface of a cone.

ence of angular momentum. This same phenomenon is responsible for the wobbling of a gyroscope, top, or *dreidel.* Because such devices have "spin," when they interact with the earth's gravitational field, they pursue the same conical trajectory.

Now let's look at the same phenomenon as it pertains to nuclei. Nuclei have charge that confers nuclear magnetism. When placed in an externally applied magnetic field, they will interact with that field in much the same way as a compass needle interacts with the earth's magnetic field: they will oscillate. If the nucleus also has spin angular momentum, it will not oscillate but will *precess* around the externally applied magnetic field, pursuing the same type of conical trajectory described above (Fig. 1-3).

What Is Spin?

The preceding description of spin angular momentum drew on examples from classical mechanics. These examples are useful but are ultimately inadequate for accurately describing the behavior of very small particles such as nuclei. Quantum mechanics is required to fully explain the effects of spin angular momentum at the nuclear level. Although a full quantum mechanical description is beyond our scope and is *not* necessary for an understanding of MRI, it is worthwhile to be aware of the quantum definition of "spin" described in the box in this section. Note that nuclear spin does not imply that the nuclei actually spin in the physical sense; in the world of quantum mechanics, small particles such as nuclei exhibit angular momentum, *but do not actually spin.*

Based on the background we have developed thus far, the "requirements" for the NMR effect can be succinctly stated as follows:

1. Nonzero charge provides nuclear magnetism.
2. Nonzero spin provides spin angular momentum.

Both of these characteristics are central in determining the nature of the interactions of nuclei with an externally applied magnetic field, and both must be present before an MRI signal or image can be obtained.

Nuclear Spin

Is a quantum mechanical phenomenon describing the behavior of very small particles;

Is proportional to spin angular momentum:

- More spin = greater spin angular momentum;

Is described by the spin number: either 0, a whole integer, or a noninteger multiple of ½:

- Spin = 0, if there are an even number of *both* protons and neutrons. Such nuclei (like He) do not exhibit NMR;
- Spin is an integer if there is an odd number of *both* protons and neutrons;
- Spin is a noninteger multiple of ½ for all other nuclei;
- Spin ranges from 0 to 7; *higher spin number = greater spin angular momentum.*

Spin determines the *gyromagnetic ratio.* This constant value is unique for each element and independent of magnetic field strength. It describes the strength of the NMR response (the frequency of precession ω) of a nucleus, and, thus, we will see it directs how we will be able to detect that nucleus using NMR.

Interaction of Protons with a Static Magnetic Field (B)

In the absence of a magnetic field external to the one exhibited by the spins, the orientation of the spins' magnetization will be completely random such that the vector sum of their magnetization will be zero. We must consider two consequences of the spins' interaction with B. First, just as a compass needle will align with the earth's magnetic field, the proton's magnetic field will align with B. Spins tend to align parallel *or* antiparallel (parallel, but pointing in the opposite direction) to B. After the spins' magnetizations align with B, the vector sum will be nonzero, yielding a net magnetic field. In this state, we say that the sample has become magnetized. The alignment of spin magnetization is in some way analogous to the preferred alignment of two permanent magnets. The "antiparallel" orientation is in fact the most likely to occur. When you press two small bar magnets together, the north pole of each will tend to align with the south pole of the other. Whereas two bar magnets will *always* assume this "antiparallel" alignment, the magnetization of our spins will align both antiparallel and parallel. This is because of the quantum mechanical nature of nuclear magnetism. The majority of the magnetization will yield a vector sum of zero with only a small excess vector sum detectable in the antiparallel orientation. This small excess is approximately equivalent to the magnetization of six nuclei in a sample of 10,000. That tiny net magnetization is the signal we must detect in order to make an MR image. However, we will soon see that this tiny component of magnetization cannot be detected unless we manipulate its orientation using the NMR effect. Second, because the spins precess about B they never fully assume a truly parallel orientation with respect to B. The influence of spin angular momentum leads to the proton's magnetization vector describing a cone (Fig. 1-3). This pattern of motion is termed *precession*. The gyromagnetic ratio—indicated by the Greek letter γ and expressed in units of MHz/T—is the unique value for each nucleus that allows us to determine the frequency (think of it as rate or speed) at which the nucleus will precess around a static magnetic field B. The gyromagnetic ratio in combination with the strength of B tells us the frequency of precession (i.e., the rate or "speed" at which each proton precesses around the axis of B), indicated by the Greek letter ω: $\omega = \gamma B$. This is the famous Larmor equation.

In the ensuing discussions, we will, until further notice, assume a constant and perfectly homogeneous externally applied magnetic field (B), referred to from now on as B_0. This means that the strength and orientation of B_0 is exactly the same in every location, so that each and every spin "sees" the same B_0. Because we will be dealing with a homogeneous magnetic field B_0, we will speak in terms of the Larmor frequency of the nucleus at that field strength. It is termed ω_0. Thus,

$$\omega_0 = \gamma B_0$$

You *must* know this equation, but this is the *only* equation that you will have to know.

Describing a Real-Life Sample: Dealing with Many Spins

In reality, whether our sample is a solution in a test tube or the human body, we never observe a single proton. Rather, billions of protons are observed *as a group.* The signal that we measure is the aggregate signal generated by *all* spins in the sample. Just as we measure the aggregate signal, in our discussions of NMR and MRI we will deal with this composite signal rather than discuss the behavior of individual spins. Although each proton is an individual, so to speak, we can describe this entire group of spins without misrepresenting the contribution of even one individual spin. The net magnetization vector (NMV), also called M_0, gives us this "bird's-eye view" of the whole sample of spins. Here we will put vector addition into practice.

In order to find the NMV, we first decompose the vector of each spin into components parallel and perpendicular to B_0. Adding the component vectors parallel to $B_{0,}$ which we call longitudinal component vectors, yields a single vector that is also parallel to B_0. This vector will be called M_z. Next we will deal with the component vectors perpendicular to B_0, calling them transverse component vectors, or M_t. Although all spins precess at the same rate (ω_0)—remember, they all "feel" the same B_0, and $\omega = \gamma B_0$—they are not all at the same point in their precession about B_0. That is, the transverse components do not all point in the same direction at any one point in time. In fact, the arrangement of the spins along the path of precession is completely random. We call this *random phase*. Given a large sample of spins—remember that we are dealing with millions, billions, or trillions of spins—we will find a transverse component vector with orientation exactly opposite each and every component vector. These transverse component vectors yield a vector sum of zero. In other words, for each transverse component vector, there is another one pointing in the opposite direction to cancel it. Recall that the magnitude of each of these component vectors is identical because each derives from a single proton experiencing an identical B_0. Random phase means that all transverse components cancel each other.

The *net* magnetization is now *correctly* represented as a single vector parallel to B_0 with a magnitude equal to the sum of the longitudinal component vectors. Note that whereas each protons' vector demonstrates precession, the NMV is stationary because the "precessing component" of magnetization has summed to zero.

The Energy Configuration Approach: A Painless (Really!) Bit of Quantum Mechanics

Overview

Explanations of the NMR phenomenon that employ principles of classical mechanics and electromagnetism abound in introductory texts on NMR and

> **A note to the faint-hearted:** This is a nonmathematical explanation of a nonetheless complex topic. It is, I believe, a more satisfying explanation than the forced classical physics descriptions. Once we understand this approach to NMR, the inconsistencies inherent in the prevalent classical explanations will be glaring. You *do not*, however, need to do the math in order to grasp the key concepts presented here.

MRI. This approach, however, inevitably runs aground, unable to really explain what is going on. This is because classical physics is not capable of explaining the behavior of very small particles such as nuclei undergoing NMR. In order to understand what is going on during NMR, we must view the events through the lens of quantum mechanics. The problem is that a full understanding of quantum mechanics is well beyond the scope of this book and beyond what most MRI users (even MRI scientists) want and need to know. So, on the one hand we have relatively easy to understand phenomena (classical approach) that do not accurately explain NMR. On the other hand, esoteric and more difficult to understand phenomena (quantum mechanics) can elegantly explain the physical basis of NMR. It is indeed tempting to take the easy path and use the classical explanations, hoping that some hand waving will suffice where the approach fails.

Another approach for those who reject the simpler but inadequate classical physics explanations is to simply accept various tenets of the NMR phenomena as postulates or articles of faith.

Our approach, on the other hand, will be to take the "high road" and invoke quantum mechanical explanations to expose the truth of MRI. The catch is that we propose to do so without delving into the more esoteric details of quantum mechanics and certainly not into the mathematics. Although the mathematics provides the clearest and most complete delineation of these concepts, it is not necessary in order to gain adequate insight to support a very good working knowledge of MRI. Our rationale for this approach is that it affords "an explanation that actually makes sense" and allows the student of MRI to see a logic behind the apparent NMR voodoo. This understanding of basic mechanisms is not an end in itself but in turn enhances understanding of concepts such as spin relaxation and facilitates retention of important concepts and their application to other aspects of NMR and MRI.

Probability and Certainty

The world of quantum mechanics is very different from the world that we are accustomed to. First, we perceive a continuous spectrum of possibilities. Human vision, for example, does not generate images composed of individual pixels but a continuous spectrum of color and objects. Second, we are accustomed to having "certain" knowledge of the state of affairs at any moment in time. In the quantum mechanical world, things cannot be "anywhere" but only in specific discrete ("quantized") states—hence the name *quantum*. Perhaps most curiously, items subject to the rules of quantum mechanics can exist in a combination of two or more different states at once. One example of this curiosity is wave-particle duality; small discrete particles can exist and behave simultaneously as continuous waves—yes, this has been shown to be true by experiment. Finally, we can never know the exact state of individual components of the system (e.g., spins), only the likely state of the system as a whole at a specific point in time. Although the quantum mechanical approach may seem unusual, it is very well-suited to describe the behavior of very small particles like hydrogen nuclei. In fact, quantum mechanics is the only truthful and satisfying means through which to understand MRI. The contrived classical explanations simply do not cut it.

Like any system, quantum mechanics operates by a set of rules. However, these rules are unlike others, such as the rules of English grammar or etiquette, in that the rules of quantum mechanics are true and absolute; there are *no* exceptions. In the world of quantum mechanics, we can describe particles (such as nuclei or protons) in terms of their location, velocity, energy, and so forth, but there are limits set on how much we can know about any one particle at any one time.

Energy Levels

Quantum mechanics identifies a discrete number of states for a given nucleus. These states are known as energy levels, and the number of energy levels varies with the type of nucleus. In the case of ^{1}H, there are only two energy levels, making it a relatively simple system to discuss. Rather than thinking of the energy levels as places where the spins physically reside, think of energy level "filling" and the relative predominance of one energy level or the other as a means to describe the overall "energy configuration" of the system.

Note that, based on the foregoing, it is *not* strictly correct to say that spins exist within or "occupy" one energy level or the other, because the energy levels describe the energy of spins, *not* their physical location. As discussed above, items subject to the rules of quantum mechanics, such as our hydrogen nuclei, can, in fact, exist in more than one state at once, termed *superposition*. The only restriction on our hydrogen protons is that they show a linear combination of the available energy levels. It is, therefore, inaccurate to describe the spins as moving between energy states. Although spins can exist with a combination of energy levels, however, there is a catch. The nature of quantum mechanics only permits us to observe spins within one energy level at a single point in time. As strange as it seems, it is actually our observation of the system that "forces" it to "choose" a single energy level. This is just another peculiar aspect of the quantized nature of the world according to quantum mechanics.

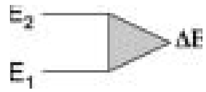

FIGURE 1-4. Energy states. Protons exist in a system of two energy states, where a shift from the lower energy state E_1 to the higher energy state E_2 requires input of an amount of energy ΔE. Conversely, a shift from E_2 to E_1 results in the release (emission) of an equal amount of energy.

Minimizing the Energy Configuration

The example shown in Figure 1-4 depicts a system with two energy states—E_1 and E_2—which is the actual number of states available to ^1H nuclei. As described above, spins must exist with energy corresponding to some linear combination of the available energy levels. When spins change their energy distribution among the available energy levels, we will observe a change in the overall energy configuration of the system. Because energy must always be conserved, in order for such a change in the energy configuration to occur, the spins must either absorb or emit energy (ΔE in Fig. 1-4), depending on the direction of shift in the energy state balance. When the lower energy level E_2 is dominant, we describe the system as being in a low-energy configuration. This is in comparison with a situation where the higher energy level E_1 predominates, yielding a high-energy configuration. The second law of thermodynamics tells us that any system will always tend toward the lowest possible energy configuration. At this lowest energy configuration, entropy (randomness) is maximal. In our case, this means an energy configuration where the lower energy state E_2 predominates. The lowest net energy configuration is called the state of maximum entropy. The corollary to this law is that maintaining a low-entropy/high-energy state requires input of energy.

These quantum mechanical laws of energy levels and distribution are ubiquitous and also govern the behavior of many things well within our daily experience. To put this into perspective, I like to think of the "household definition of entropy": Your home is a system, and, if it is to function according to the second law of thermodynamics, there must be a distribution among energy states. Organization, where each of your possessions is in its place, represents a high-energy configuration, and complete disarray is a low-energy configuration. The second law of thermodynamics tells us that things will tend to become disorganized; to maintain order at home, we must put energy into the system (home) in the form of cleaning effort.

If any system will move to maximize entropy by shifting to the lowest energy configuration possible, why is the higher energy level E_1 not vacated entirely? This is because of thermal energy within the spin system. At absolute zero, where molecular motion stops because of the complete absence of thermal energy, spins will in fact only exist at the lowest energy level.

The Uncertainty Principle

At all temperatures exceeding absolute zero (see earlier discussion), the distribution of spins among the available energy levels is not static but constantly in flux. The Heisenberg uncertainty principle dictates that we cannot know the state of an individual spin at a discrete point in time. Nonetheless, we can know the probability of finding a certain distribution of spins among the available energy states, given a specific set of ambient conditions. We can demonstrate this by making experimental measurements of the system. We will always be able to confirm our predictions about the energy configuration *as a* whole by making a large number of measurements.

The corollary of the description above follows: by measuring the energy configuration of the system (something that we will measure as the MRI signal intensity), we will be able to infer the distribution of spins among energy states. Once again, these statements apply *only* to the system as a whole, not to the specific state of individual spins. In the following section, we will see how the MRI signal intensity can be most completely described in terms of these energy distributions.

Summary

These, then, are the ground rules for our quantum mechanical approach:

- We can never know the energy state or location of any individual spin, only the *probability* of finding spins in a certain distribution among energy states under given conditions. We are only permitted to talk about the configuration of the whole system.

> *What any one spin is doing is none of your business.*

- Energy levels describe the nature of the system but do not imply physical location of the spins in one level or another. Spins' energy can exist with a combination of energy levels.
- The predominance of specific energy levels can be described as the energy configuration of the whole system.
- The system energy configuration cannot change without input or emission of energy.

Making the Quantum Mechanical Approach Specific to MRI

First, let us identify the energy configurations we will be dealing with. As we have already established, all of our discussions here will concern the NMV and its longitudinal and transverse components. Earlier when we introduced the NMV, it had *no* transverse component because all of the individual spins that contribute to the NMV have random phase. As a result, their transverse components cancel out. This is true if we place our sample in the B_0 field and observe it at rest. However, we will see shortly that, in order to measure the MRI signal, we must disturb this resting state and create a net transverse component to the NMV.

In a low-energy configuration where population of the lower energy level (E_1) predominates, we will observe the longitudinal component of magnetization—we will call this M_z from now on—parallel to B_0. In addition, in this low-energy configuration, we will observe that random phase predominates. This means that there will be no net transverse magnetization (referred to as M_t from now on); all the randomly oriented transverse components will add to 0. Note that this is the resting state where we observe the NMV as originally described: a stationary vector parallel to B_0.

A higher-energy configuration, where population of the higher energy level (E_2) is greater than at rest, can only be observed if energy is added to the system. The details of how we add energy to disturb the resting NMV and what happens when we disturb it are the subject of the next chapter. When we do observe this higher-energy configuration, net longitudinal magnetization (M_z) will be lower than at rest. Simultaneously, phase coherence of the spins contributing to the NMV will increase, yielding net transverse magnetization (M_t). The degree to which we observe these phenomena—decreasing M_z and increasing M_t—depends on the extent to which we shift toward a higher-energy configuration. This, of course, is a function of the amount of energy that we add to the system.

At this point it should be very clear how essential vector addition is to understanding the behavior of spins in a magnetic field. (See Appendix 1 if vectors and their addition and decomposition are not completely clear.) As M_z declines and M_t increases, the vector sum (i.e., the NMV—originally parallel to B_0) does not change in magnitude but becomes oriented at an angle to B_0. If enough energy is deposited, the NMV will form a 90° angle with respect to B_0 (Fig. 1-5), and, if even more energy is deposited, the orientation with respect to B_0 will exceed 90°. How we achieve this "flip" in the orientation of the NMV with respect to B_0 is, believe it or not, *best* understood through the distribution of spins among energy states and how that distribution is altered by the input of energy. Attempts at classical (mechanical/electromagnetism) explanations of this phenomenon simply fall flat on their face.

Remember that energy must always be conserved. Just as adding energy to the system will cause a shift of spins toward a higher-energy configuration, redistribution of spins back to the lower-energy state requires loss of energy from the spin system. In a subsequent chapter, we will discuss how energy can be transferred to the environment or to other spins to accomplish this return to a lower (resting) configuration. Finally, notice that the longitudinal and transverse components are in balance:

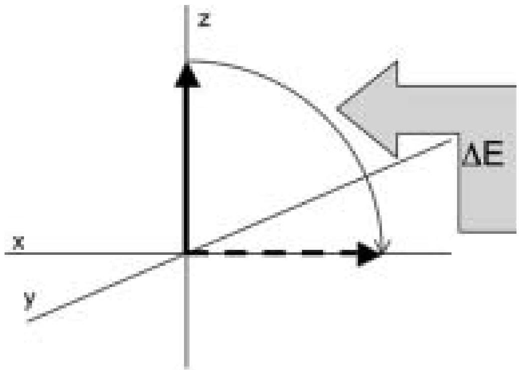

FIGURE 1-5. Excitation. In this example, B_0 is parallel to the z axis. The consequence of addition of energy and the resultant shift to a higher-energy configuration is the rotation of the NMV from its alignment parallel to B_0 (thick arrow) toward one perpendicular to it (broken arrow).

- Adding energy yields: $\Downarrow M_z$ *and* $\Uparrow M_t$ = flip of NMV of 90°, or more or less depending on the amount of energy added;
- Loss of energy yields: $\Uparrow M_z$ *and* $\Downarrow M_t$ = return of NMV to rest (M_0) parallel to B_0.

The following chapters will address how we cause these shifts in the energy state of the system and will examine in much greater detail what happens when we sit back and watch things move back to "normal": maximum entropy/lowest energy configuration, that is, of course!

One More Thing . . . What Exactly Is the MRI Signal That We Measure?

As we will see in the next chapter, MR signal intensity is the measured magnitude of M_t. If we are measuring magnetization, why not just measure M_z? The short answer is that the MRI signal from tissue is so small relative to B_0 that *we would be looking for a needle in the proverbial haystack.* If we chose to use a gaussmeter (a device used to measure magnetic field strength) to measure the magnetization of tissue, our instrument would be swamped by the huge amount of B_0 magnetization. The tiny amount of tissue magnetization would be undetectable. In the next chapter, we will learn that, because M_t is in a different *orientation* than B_0, we can detect it independently and with great sensitivity.

http://www.springer.com/978-0-387-48895-0

Totally Accessible MRI
A User's Guide to Principles, Technology, and Applications
Lipton, M.L.
2008, XXII, 384 p. 153 illus., 7 illus. in color., Softcover
ISBN: 978-0-387-48895-0